

in 10/22  
chs 19,12

# M. E. Williams, Inc.

## Enclosure Memo

October 20, 2010

TO: Martha Wetherill

FROM: Martha Williams

RE: Greene 2140242

Dear Martha,

Please find herewith:

Copyedited manuscript for Chapters 19 and 20

Sincerely,

Martha

## LIMITED DEPENDENT VARIABLES TRUNCATION, CENSORING AND SAMPLE SELECTION

### 19.1 Introduction

This chapter is concerned with **truncation** and **censoring**. As we saw in Section 18.4.6, these features complicate the analysis of data that might otherwise be amenable to conventional estimation methods such as regression. "Truncation" effects arise when one attempts to make inferences about a larger population from a sample that is drawn from a distinct subpopulation. For example, studies of income based on incomes above or below some poverty line may be of limited usefulness for inference about the whole population. Truncation is essentially a characteristic of the distribution from which the sample data are drawn. Censoring is a more common feature of recent studies. To continue the example, suppose that instead of being unobserved, all incomes below the poverty line are reported as if they were *at* the poverty line. The censoring of a range of values of the variable of interest introduces a distortion into conventional statistical results that is similar to that of truncation. Unlike truncation, however, censoring is essentially a defect in the sample data. Presumably, if they were not censored, the data would be a representative sample from the population of interest. We will also examine a form of truncation called the **sample selection** problem. Although most empirical work in this area involves censoring rather than truncation, we will study the simpler model of truncation first. It provides most of the theoretical tools we need to analyze models of censoring and sample selection.

✓ and/ The discussion will examine the general characteristics of truncation, censoring and sample selection, then, in each case, develop a major area of application of the principles. The stochastic frontier model [Aigner, Lovell, and Schmidt (1977), Fried, Lovell and Schmidt (2008)] is a leading application of results for truncated distributions in empirical models. Censoring appears prominently in the analysis of labor supply and in modeling of duration data. Finally, the sample selection model has appeared in all areas of the social sciences, and plays a significant role in the evaluation of treatment effects and program evaluation.

18 19  
24

# TRUNCATION, CENSORING, AND SAMPLE SELECTION

## 24.1 INTRODUCTION

This chapter is concerned with truncation and censoring.<sup>1</sup> The effect of truncation occurs when sample data are drawn from a subset of a larger population of interest. For example, studies of income based on incomes above or below some poverty line may be of limited usefulness for inference about the whole population. Truncation is essentially a characteristic of the distribution from which the sample data are drawn. Censoring is a more common feature of recent studies. To continue the example, suppose that instead of being unobserved, all incomes below the poverty line are reported as if they were at the poverty line. The censoring of a range of values of the variable of interest introduces a distortion into conventional statistical results that is similar to that of truncation. Unlike truncation, however, censoring is essentially a defect in the sample data. Presumably, if they were not censored, the data would be a representative sample from the population of interest.

This chapter will discuss three broad topics: truncation, censoring, and a form of truncation called the **sample selection** problem. Although most empirical work in this area involves censoring rather than truncation, we will study the simpler model of truncation first. It provides most of the theoretical tools we need to analyze models of censoring and sample selection.<sup>2</sup>

## 19.2 ~~19.2.1~~ TRUNCATION

In this section, we are concerned with inferring the characteristics of a full population from a sample drawn from a restricted part of that population.

### 19.2.1 ~~24.2.1~~ TRUNCATED DISTRIBUTIONS

19.2

A **truncated distribution** is the part of an untruncated distribution that is above or below some specified value. For instance, in Example ~~24.2~~, we are given a characteristic of the distribution of incomes above \$100,000. This subset is a part of the full distribution of incomes which range from zero to (essentially) infinity.

<sup>1</sup> Five of the many surveys of these topics are Dhrymes (1984), Maddala (1977b, 1983, 1984) and Amemiya (1984). The last is part of a symposium on censored and truncated regression models. Surveys that are oriented toward applications and techniques are Long (1997), deMaris (2004), and Greene (2006). Some recent results on non- and semiparametric estimation appear in Lee (1996).

<sup>2</sup> Although sample selection can be viewed formally as merely a type of truncation, it is useful to treat the methodological aspects separately.

FN 1

**19.1 THEOREM 24.1** Density of a Truncated Random Variable  
 If a continuous random variable  $x$  has pdf  $f(x)$  and  $a$  is a constant, then

$$f(x|x > a) = \frac{f(x)}{\text{Prob}(x > a)}$$

The proof follows from the definition of conditional probability and amounts merely to scaling the density so that it integrates to one over the range above  $a$ . Note that the truncated distribution is a conditional distribution.

Most recent applications based on continuous random variables use the **truncated normal distribution**. If  $x$  has a normal distribution with mean  $\mu$  and standard deviation  $\sigma$ , then

$$\text{Prob}(x > a) = 1 - \Phi\left(\frac{a - \mu}{\sigma}\right) = 1 - \Phi(\alpha),$$

where  $\alpha = (a - \mu)/\sigma$  and  $\Phi(\cdot)$  is the standard normal cdf. The density of the truncated normal distribution is then

$$f(x|x > a) = \frac{f(x)}{1 - \Phi(\alpha)} = \frac{(2\pi\sigma^2)^{-1/2} e^{-(x-\mu)^2/(2\sigma^2)}}{1 - \Phi(\alpha)} = \frac{\frac{1}{\sigma} \phi\left(\frac{x-\mu}{\sigma}\right)}{1 - \Phi(\alpha)}$$

FIG 19.1

where  $\phi(\cdot)$  is the standard normal pdf. The **truncated standard normal distribution**, with  $\mu = 0$  and  $\sigma = 1$ , is illustrated for  $a = -0.5, 0$ , and  $0.5$  in Figure 24.1. Another truncated distribution that has appeared in the recent literature, this one for a discrete random variable, is the truncated at zero Poisson distribution.

$$\begin{aligned} \text{Prob}[Y = y | y > 0] &= \frac{(e^{-\lambda} \lambda^y)/y!}{\text{Prob}[Y > 0]} = \frac{(e^{-\lambda} \lambda^y)/y!}{1 - \text{Prob}[Y = 0]} \\ &= \frac{(e^{-\lambda} \lambda^y)/y!}{1 - e^{-\lambda}}, \quad \lambda > 0, y = 1, \dots \end{aligned}$$

FN 2

This distribution is used in models of uses of recreation and other kinds of facilities where observations of zero uses are discarded.

For convenience in what follows, we shall call a random variable whose distribution is truncated a **truncated random variable**.

19.2.2 ~~24.2.2~~ MOMENTS OF TRUNCATED DISTRIBUTIONS

We are usually interested in the mean and variance of the truncated random variable. They would be obtained by the general formula:

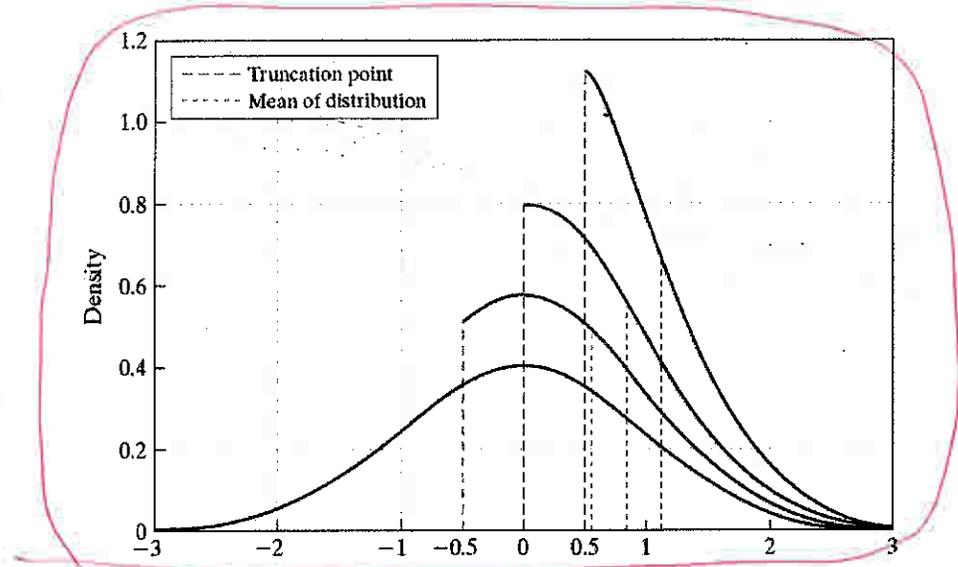
$$E[x|x > a] = \int_a^\infty x f(x|x > a) dx$$

for the mean and likewise for the variance.

<sup>3</sup>The case of truncation from above instead of below is handled in an analogous fashion and does not require any new results.

<sup>4</sup>See Shaw (1988). An application of this model appears in Section 18.4.6 and Example 18.8

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 865



19.1  
**FIGURE 24.1** Truncated Normal Distributions.

19.1  
**Example 24.1** *Truncated Uniform Distribution*  
If  $x$  has a standard uniform distribution, denoted  $U(0, 1)$ , then

$$f(x) = 1, \quad 0 \leq x \leq 1.$$

The truncated at  $x = \frac{1}{3}$  distribution is also uniform;

$$f\left(x \mid x > \frac{1}{3}\right) = \frac{f(x)}{\text{Prob}\left(x > \frac{1}{3}\right)} = \frac{1}{\left(\frac{2}{3}\right)} = \frac{3}{2}, \quad \frac{1}{3} \leq x \leq 1.$$

The expected value is

$$E\left[x \mid x > \frac{1}{3}\right] = \int_{1/3}^1 x \left(\frac{3}{2}\right) dx = \frac{2}{3}.$$

For a variable distributed uniformly between  $L$  and  $U$ , the variance is  $(U - L)^2/12$ . Thus,

$$\text{Var}\left[x \mid x > \frac{1}{3}\right] = \frac{1}{27}.$$

The mean and variance of the untruncated distribution are  $\frac{1}{2}$  and  $\frac{1}{12}$ , respectively.

19  
Example 24.1 illustrates two results.

1. If the truncation is from below, then the mean of the truncated variable is greater than the mean of the original one. If the truncation is from above, then the mean of the truncated variable is smaller than the mean of the original one. ~~This is clearly visible in Figure 24.1.~~
2. Truncation reduces the variance compared with the variance in the untruncated distribution.

Henceforth, we shall use the terms truncated mean and truncated variance to refer to the mean and variance of the random variable with a truncated distribution.

19-5

866 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

FN  
3

For the truncated normal distribution, we have the following theorem:<sup>5</sup>

**THEOREM 19.2** ~~24.2~~ Moments of the Truncated Normal Distribution  
 If  $x \sim N[\mu, \sigma^2]$  and  $a$  is a constant, then

$$E[x | \text{truncation}] = \mu + \sigma \lambda(\alpha), \quad (19-24-1)$$

$$\text{Var}[x | \text{truncation}] = \sigma^2 [1 - \delta(\alpha)], \quad (19-24-2)$$

where  $\alpha = (a - \mu)/\sigma$ ,  $\phi(\alpha)$  is the standard normal density and

$$\lambda(\alpha) = \phi(\alpha) / [1 - \Phi(\alpha)] \quad \text{if truncation is } x > a, \quad (19-24-3a)$$

$$\lambda(\alpha) = -\phi(\alpha) / \Phi(\alpha) \quad \text{if truncation is } x < a, \quad (19-24-3b)$$

and

$$\delta(\alpha) = \lambda(\alpha) [\lambda(\alpha) - \alpha]. \quad (19-24-4)$$

An important result is

$$0 < \delta(\alpha) < 1 \quad \text{for all values of } \alpha,$$

which implies point 2 after Example 19.1. A result that we will use at several points below is  $d\phi(\alpha)/d\alpha = -\alpha\phi(\alpha)$ . The function  $\lambda(\alpha)$  is called the **inverse Mills ratio**. The function in (24-3a) is also called the **hazard function** for the standard normal distribution. (KT)

19c

**Example 19.2** A Truncated Lognormal Income Distribution

"The typical 'upper affluent American' ... makes \$142,000 per year. ... The people surveyed had household income of at least \$100,000."<sup>6</sup> Would this statistic tell us anything about the "typical American"? As it stands, it probably does not (popular impressions notwithstanding). The 1987 article where this appeared went on to state, "If you're in that category, pat yourself on the back—only 2 percent of American households make the grade, according to the survey." Because the degree of truncation in the sample is 98 percent, the \$142,000 was probably quite far from the mean in the full population.

4

FN  
H

(KT)

Suppose that incomes,  $x$ , in the population were lognormally distributed—see Section B.4.4. Then, the log of income,  $y$ , had a normal distribution with, say, mean  $\mu$  and standard deviation,  $\sigma$ . Suppose that the survey was large enough for us to treat the sample average as the true mean. Assuming so, we'll deduce  $\mu$  and  $\sigma$  then determine the population mean income. and/

Two useful numbers for this example are  $\ln 100 = 4.605$  and  $\ln 142 = 4.956$ . The article states that

$$\text{Prob}[x \geq 100] = \text{Prob}[\exp(y) \geq 100] = 0.02,$$

or

$$\text{Prob}(y < 4.605) = 0.98.$$

✓  
✓

3  
4

<sup>5</sup>Details may be found in Johnson, Kotz, and Balakrishnan (1994, pp. 156–158). Prof's appear in

<sup>6</sup>See *New York Post* (1987).

Cameron and Trivedi; (2005).

## CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 867

This implies that

$$\text{Prob}[(y - \mu)/\sigma < (4.605 - \mu)/\sigma] = 0.98.$$

Because  $\Phi[(4.605 - \mu)/\sigma] = 0.98$ , we know that

$$\Phi^{-1}(0.98) = 2.054 = (4.605 - \mu)/\sigma,$$

or

$$4.605 = \mu + 2.054\sigma.$$

The article also states that

$$E[x | x > 100] = E[\exp(y) | \exp(y) > 100] = 142,$$

or

$$E[\exp(y) | y > 4.645] = 142.$$

To proceed, we need another result for the lognormal distribution:

$$\text{If } y \sim N[\mu, \sigma^2], \text{ then } E[\exp(y) | y > a] = \exp(\mu + \sigma^2/2) \times \frac{\Phi(\sigma - (a - \mu)/\sigma)}{1 - \Phi((a - \mu)/\sigma)}.$$

[See Johnson, Kotz and Balakrishnan (1995, p. 241).] For our application, we would equate this expression to 142, and  $a$  to  $\ln 100 = 4.605$ . This provides a second equation. To estimate the two parameters, we used the method of moments. We solved the minimization problem

$$\text{Minimize}_{\mu, \sigma} [4.605 - (\mu + 2.054\sigma)]^2 + [142\Phi((\mu - 4.605)/\sigma) - \exp(\mu + \sigma^2/2)\Phi(\sigma - (4.605 - \mu)/\sigma)]^2.$$

The two solutions are 2.89372 and 0.83314 for  $\mu$  and  $\sigma$ , respectively. To obtain the mean income, we now use the result that if  $y \sim N[\mu, \sigma^2]$  and  $x = \exp(y)$ , then  $E[x] = \exp(\mu + \sigma^2/2)$ . Inserting our values for  $\mu$  and  $\sigma$  gives  $E[x] = \$25,554$ . The 1987 Statistical Abstract of the United States gives the mean of household incomes across all groups for the United States as about \$25,000. So, the estimate based on surprisingly little information would have been relatively good. These meager data did, indeed, tell us something about the average American.

19.2.3 ~~24.2.3~~ THE TRUNCATED REGRESSION MODEL

In the model of the earlier examples, we now assume that

$$\mu_i = \mathbf{x}_i' \boldsymbol{\beta}$$

is the deterministic part of the classical regression model. Then

$$y_i = \mathbf{x}_i' \boldsymbol{\beta} + \varepsilon_i,$$

where

$$\varepsilon_i | \mathbf{x}_i \sim N[0, \sigma^2],$$

so that

$$y_i | \mathbf{x}_i \sim N[\mathbf{x}_i' \boldsymbol{\beta}, \sigma^2].$$

868 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

We are interested in the distribution of  $y_i$  given that  $y_i$  is greater than the truncation point  $a$ . This is the result described in Theorem 19.2. It follows that

$$E[y_i | y_i > a] = x_i' \beta + \sigma \frac{\phi[(a - x_i' \beta) / \sigma]}{1 - \Phi[(a - x_i' \beta) / \sigma]} \quad \begin{matrix} 19 \\ (24-6) \end{matrix}$$

The conditional mean is therefore a nonlinear function of  $a$ ,  $\sigma$ ,  $x$ , and  $\beta$ .

The ~~marginal~~ <sup>partial</sup> effects in this model *in the subpopulation* can be obtained by writing

$$E[y_i | y_i > a] = x_i' \beta + \sigma \lambda(\alpha_i), \quad \begin{matrix} 19 \\ (24-7) \end{matrix}$$

where now  $\alpha_i = (a - x_i' \beta) / \sigma$ . For convenience, let  $\lambda_i = \lambda(\alpha_i)$  and  $\delta_i = \delta(\alpha_i)$ . Then

$$\begin{aligned} \frac{\partial E[y_i | y_i > a]}{\partial x_i} &= \beta + \sigma (d\lambda_i / d\alpha_i) \frac{\partial \alpha_i}{\partial x_i} \\ &= \beta + \sigma (\lambda_i^2 - \alpha_i \lambda_i) (-\beta / \sigma) \\ &= \beta (1 - \lambda_i^2 + \alpha_i \lambda_i) \\ &= \beta (1 - \delta_i). \end{aligned} \quad \begin{matrix} 19 \\ (24-8) \end{matrix}$$

Note the appearance of the scale factor  $1 - \delta_i$  from the truncated variance. Because  $(1 - \delta_i)$  is between zero and one, we conclude that for every element of  $x_i$ , the marginal effect is less than the corresponding coefficient. There is a similar *attenuation* of the variance. In the subpopulation  $y_i > a$ , the regression variance is not  $\sigma^2$  but

$$\text{Var}[y_i | y_i > a] = \sigma^2 (1 - \delta_i). \quad \begin{matrix} 19 \\ (24-9) \end{matrix}$$

Whether the ~~marginal~~ <sup>partial</sup> effect in (24-7) or the coefficient  $\beta$  itself is of interest depends on the intended inferences of the study. If the analysis is to be confined to the subpopulation, then (24-7) is of interest. If the study is intended to extend to the entire population, however, then it is the coefficients  $\beta$  that are actually of interest.

One's first inclination might be to use ordinary least squares to estimate the parameters of this regression model. For the subpopulation from which the data are drawn, we could write (24-6) in the form

$$y_i | y_i > a = E[y_i | y_i > a] + u_i = x_i' \beta + \sigma \lambda_i + u_i, \quad \begin{matrix} 19 \\ (24-10) \end{matrix}$$

where  $u_i$  is  $y_i$  minus its conditional expectation. By construction,  $u_i$  has a zero mean, but it is heteroscedastic:

$$\text{Var}[u_i] = \sigma^2 (1 - \lambda_i^2 + \lambda_i \alpha_i) = \sigma^2 (1 - \delta_i),$$

which is a function of  $x_i$ . If we estimate (24-10) by ordinary least squares regression of  $y$  on  $X$ , then we have omitted a variable, the nonlinear term  $\lambda_i$ . All the biases that arise because of an omitted variable can be expected.

Without some knowledge of the distribution of  $x$ , it is not possible to determine how serious the bias is likely to be. A result obtained by Chung and Goldberger (1984) is broadly suggestive. If  $E[x | y]$  in the full population is a linear function of  $y$ , then  $\text{plim } \hat{b} = \beta \tau$  for some proportionality constant  $\tau$ . This result is consistent with the widely observed (albeit rather rough) proportionality relationship between least squares estimates

<sup>7</sup>See Heckman (1979) who formulates this as a "specification error."

FN 5

## CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 869

of this model and maximum likelihood estimates.<sup>8</sup> The proportionality result appears to be quite general. In applications, it is usually found that, compared with consistent maximum likelihood estimates, the OLS estimates are biased toward zero. (See Example 19.5.)

## 24.3 CENSORED DATA

A very common problem in microeconomic data is **censoring** of the dependent variable. When the dependent variable is censored, values in a certain range are all transformed to (or reported as) a single value. Some examples that have appeared in the empirical literature are as follows:<sup>9</sup>

1. Household purchases of durable goods [Tobin (1958)].
2. The number of extramarital affairs [Fair (1977, 1978)].
3. The number of hours worked by a woman in the labor force [Quester and Greene (1982)].
4. The number of arrests after release from prison [Witte (1980)].
5. Household expenditure on various commodity groups [Jarque (1987)].
6. Vacation expenditures [Melenberg and van Soest (1996)].

Each of these studies analyzes a dependent variable that is zero for a significant fraction of the observations. Conventional regression methods fail to account for the qualitative difference between *limit* (zero) observations and *nonlimit* (continuous) observations.

## 24.3.1 THE CENSORED NORMAL DISTRIBUTION

The relevant distribution theory for a **censored variable** is similar to that for a truncated one. Once again, we begin with the normal distribution, as much of the received work has been based on an assumption of normality. We also assume that the censoring point is zero, although this is only a convenient normalization. In a truncated distribution, only the part of distribution above  $y = 0$  is relevant to our computations. To make the distribution integrate to one, we scale it up by the probability that an observation in the untruncated population falls in the range that interests us. When data are censored, the distribution that applies to the sample data is a mixture of discrete and continuous distributions. Figure 24.2 illustrates the effects.

To analyze this distribution, we define a new random variable  $y$  transformed from the original one,  $y^*$ , by

$$y = 0 \quad \text{if } y^* \leq 0,$$

$$y = y^* \quad \text{if } y^* > 0.$$

The distribution that applies if  $y^* \sim N[\mu, \sigma^2]$  is  $\text{Prob}(y=0) = \text{Prob}(y^* \leq 0) = \Phi(-\mu/\sigma) = 1 - \Phi(\mu/\sigma)$ , and if  $y^* > 0$ , then  $y$  has the density of  $y^*$ .

This distribution is a mixture of discrete and continuous parts. The total probability is one, as required, but instead of scaling the second part, we simply assign the full probability in the censored region to the censoring point, in this case, zero.

<sup>8</sup>See the appendix in Hausman and Wise (1977) and Greene (1983) as well.

<sup>9</sup>More extensive listings may be found in Amemiya (1984) and Maddala (1983).

# 19.2.4 The Stochastic Frontier Model

~~588 PART IV Estimation Methodology~~

likelihood is used when the disturbances in a regression, or the dependent variable, more generally, is not normally distributed. We now consider two applications.

19.2.4

~~19.0.5.a Nonnormal Disturbances, The Stochastic Frontier Model~~

~~This application will examine a regressionlike model in which the disturbances do not have a normal distribution. The model developed here also presents a convenient platform on which to illustrate the use of the invariance property of maximum likelihood estimators to simplify the estimation of the model.~~

A lengthy literature commencing with theoretical work by Knight (1933), Debreu (1951), and Farrell (1957) and the pioneering empirical study by Aigner, Lovell, and Schmidt (1977) has been directed at models of production that specifically account for the textbook proposition that a production function is a theoretical ideal.<sup>2</sup> If  $y = f(x)$  defines a production relationship between inputs,  $x$ , and an output,  $y$ , then for any given  $x$ , the observed value of  $y$  must be less than or equal to  $f(x)$ . The implication for an empirical regression model is that in a formulation such as  $y = h(x, \beta) + u$ ,  $u$  must be negative. Because the theoretical production function is an ideal—the frontier of efficient production—any nonzero disturbance must be interpreted as the result of inefficiency. A strictly orthodox interpretation embedded in a Cobb–Douglas production model might produce an empirical frontier production model such as

$$\ln y = \beta_1 + \sum_k \beta_k \ln x_k - u, \quad u \geq 0.$$

The gamma model described in Example 4.9 was an application. One-sided disturbances such as this one present a particularly difficult estimation problem. The primary theoretical problem is that any measurement error in  $\ln y$  must be embedded in the disturbance. The practical problem is that the entire estimated function becomes a slave to any single errantly measured data point.

Aigner, Lovell, and Schmidt proposed instead a formulation within which observed deviations from the production function could arise from two sources: (1) productive inefficiency, as we have defined it earlier and that would necessarily be negative, and (2) idiosyncratic effects that are specific to the firm and that could enter the model with either sign. The end result was what they labeled the stochastic frontier.

$$\begin{aligned} \ln y &= \beta_1 + \sum_k \beta_k \ln x_k - u + v, \quad u \geq 0, \quad v \sim N[0, \sigma_v^2] \\ &= \beta_1 + \sum_k \beta_k \ln x_k + \varepsilon. \end{aligned}$$

The frontier for any particular firm is  $h(x, \beta) + v$ , hence the name stochastic frontier. The inefficiency term is  $u$ , a random variable of particular interest in this setting. Because the data are in log terms,  $u$  is a measure of the percentage by which the particular observation fails to achieve the frontier, ideal production rate.

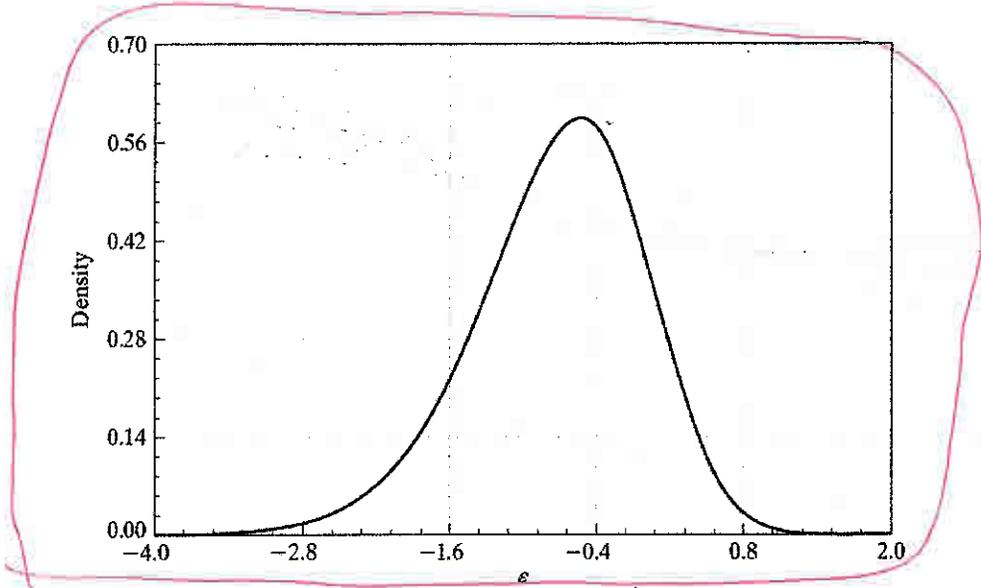
To complete the specification, they suggested two possible distributions for the inefficiency term: the absolute value of a normally distributed variable and an exponentially distributed variable. The density functions for these two compound variables are given by Aigner, Lovell, and Schmidt: let  $\varepsilon = v - u$ ,  $\lambda = \sigma_u / \sigma_v$ ,  $\sigma = (\sigma_u^2 + \sigma_v^2)^{1/2}$ ,

which has the truncated at zero distribution shown in Figure 19.1,

<sup>2</sup>A survey by Greene (2007a) appears in Fried, Lovell, and Schmidt (2007). Kumbhakar and Lovell (2000) is a comprehensive reference on the subject.

ALS, FN 7

4.7



**FIGURE 16.4** Density for the Disturbance in the Stochastic Frontier Model.  
19.2

and  $\Phi(z)$  = the probability to the left of  $z$  in the standard normal distribution (see Section B.4.1). For the “half-normal” model,

$$\ln h(\varepsilon_i | \beta, \lambda, \sigma) = \left[ -\ln \sigma + \left(\frac{1}{2}\right) \ln \frac{2}{\pi} - \frac{1}{2} \left(\frac{\varepsilon_i}{\sigma}\right)^2 + \ln \Phi\left(\frac{-\varepsilon_i \lambda}{\sigma}\right) \right],$$

whereas for the exponential model

$$\ln h(\varepsilon_i | \beta, \theta, \sigma_v) = \left[ \ln \theta + \frac{1}{2} \theta^2 \sigma_v^2 + \theta \varepsilon_i + \ln \Phi\left(-\frac{\varepsilon_i}{\sigma_v} - \theta \sigma_v\right) \right].$$

Both these distributions are asymmetric. We thus have a regression model with a nonnormal distribution specified for the disturbance. The disturbance,  $\varepsilon$ , has a nonzero mean as well:  $E[\varepsilon] = -\sigma_v(2/\pi)^{1/2}$  for the half-normal model and  $-1/\theta$  for the exponential model. Figure 16.4 illustrates the density for the half-normal model with  $\sigma = 1$  and  $\lambda = 2$ . By writing  $\beta_0 = \beta_1 + E[\varepsilon]$  and  $\varepsilon^* = \varepsilon - E[\varepsilon]$ , we obtain a more conventional formulation

$$\ln y = \beta_0 + \sum_k \beta_k \ln x_k + \varepsilon^*,$$

which does have a disturbance with a zero mean but an asymmetric, nonnormal distribution. The asymmetry of the distribution of  $\varepsilon^*$  does not negate our basic results for least squares in this classical regression model. This model satisfies the assumptions of the Gauss-Markov theorem, so least squares is unbiased and consistent (save for the constant term) and efficient among linear unbiased estimators. In this model, however, the maximum likelihood estimator is not linear, and it is more efficient than least squares.

FIG 19.2  
19.2

The log-likelihood function for the half normal model is given in ALS (1977):

$$\ln L = -n \ln \sigma + \frac{n}{2} \ln \frac{2}{\pi} - \frac{1}{2} \sum_{i=1}^n \left( \frac{\varepsilon_i}{\sigma} \right)^2 + \sum_{i=1}^n \ln \Phi \left( \frac{-\varepsilon_i \lambda}{\sigma} \right) \quad (19-11)$$

Maximization programs for this model are built into modern software packages such as *Stata*, *NLOGIT* and *TSP*. The log likelihood is simple enough that it can also be readily adapted to the generic optimization routines in, e.g., *MatLab* or *Gauss*. Some treatments in the literature use the parameterization employed by Battese and Coelli (1992) and Coelli (1996),  $\gamma = \sigma_u^2 / \sigma^2$ . This is a one to one transformation of  $\lambda$ ;  $\lambda = (\gamma / (1-\gamma))^{1/2}$ , so which parameterization is employed is a matter of convenience; the empirical results will be the same. The log likelihood function for the exponential model can be built up from the density given earlier. For the half-normal model, we would also rely on the invariance of maximum likelihood estimators to recover estimates of the structural variance parameters,  $\sigma_v^2 = \sigma^2 / (1 + \lambda^2)$  and  $\sigma_u^2 = \sigma^2 \lambda^2 / (1 + \lambda^2)$ .<sup>10</sup> (Note, the variance of the truncated variable,  $u_i$  is not  $\sigma_u^2$ ; using (19-2), it reduces to  $(1 - 2/\pi)\sigma_u^2$ .) In addition, a structural parameter of interest is the proportion of the total variance of  $\varepsilon$  that is due to the inefficiency term. For the half-normal model,  $\text{Var}[\varepsilon] = \text{Var}[u] + \text{Var}[v] = (1 - 2/\pi)\sigma_u^2 + \sigma_v^2$  whereas for the exponential model, the counterpart is  $1/\theta^2 + \sigma_v^2$ .

Modeling in the stochastic frontier setting is rather unlike what we are accustomed to up to this point, in that the disturbance, specifically  $u_i$ , not the model parameters, is the central focus of the analysis. The reason is that in this context, the disturbance,  $u_i$ , rather than being the catchall for the unknown and unknowable factors omitted from the equation, has a particular interpretation — it is the firm specific inefficiency. Ideally, we would like to estimate  $u_i$  for each firm in the sample to compare them on the basis of their productive efficiency. Unfortunately, the data do not permit a direct estimate, because with estimates of  $\beta$  in hand, we are only able to compute a direct estimate of  $\varepsilon_i = y_i - x_i' \beta$ . Jondrow et al. (1982), however, have derived a useful approximation that is now the standard measure in these settings,

$$E[u_i | \varepsilon_i] = \frac{\sigma \lambda}{1 + \lambda^2} \left[ \frac{\phi(z_i)}{1 - \Phi(z_i)} - z_i \right], \quad z_i = \frac{\varepsilon_i \lambda}{\sigma}$$

for the half-normal model, and

$$E[u_i | \varepsilon_i] = z_i + \sigma_v \frac{\phi(z_i / \sigma_v)}{\Phi(z_i / \sigma_v)}, \quad z_i = -(\varepsilon_i + \theta \sigma_v^2)$$

for the exponential model. These values can be computed using the maximum likelihood estimates of the structural parameters in the model. In some cases in which researchers are interested in discovering best practice [e.g., WHO (2000), Tandon et al. (2000)], the estimated values are sorted and the ranks of the individuals in the sample become of interest.

<sup>10</sup> A vexing problem for estimation of the model is that if the ordinary least squares residuals are skewed in the positive (wrong) direction (See Figure 19.2), OLS with  $\hat{\lambda} = 0$  will be the MLE. OLS residuals with a positive skew are apparently inconsistent with a model in which, in theory, they should have a negative skew. [See Waldman (1982) for theoretical development of this result.]

AD: OK to spell out "e.g."?

for example  
minus  
FN  
8

minus

Av: I'm  
text  
missing  
here?

Research in this arena since the methodological developments beginning in the 1930s and the building of the empirical foundations in 1977 and 1982, has proceeded in several directions. The research agenda has included the following areas:

Most theoretical treatments of "inefficiency" as envisioned here attribute it to aspects of management of the firm. It remains to establish a firm theoretical connection between the theory of firm behavior and the stochastic frontier model as a device for measurement of inefficiency.

In the context of the model, many studies have developed alternative, more flexible functional forms that (it is hoped) can provide a more realistic model for inefficiency. Two that are relevant in this chapter are Stevenson's (1980) truncated normal model and the normal-gamma frontier. One intuitively appealing form of the truncated normal model is

$$U_i \sim N[\mu + z_i'\alpha, \sigma_u^2],$$
$$u_i = |U_i|.$$

The original normal - half normal model results if  $\mu$  equals zero and  $\alpha$  equals zero. This is a device by which environmental variables noted in the next paragraph can enter the model of inefficiency. A truncated normal model is presented in Example 19.3. The half normal, truncated normal and exponential models all take the form of distribution shown in Figure 19.1. The gamma model,

$$f(u) = [\theta^P/\Gamma(P)] \exp(-\theta u) u^{P-1},$$

is a flexible model that presents the advantage that the distribution of inefficiency can move away from zero. If  $P$  is greater than one, then the density at  $u = 0$  equals zero and the entire distribution moves away from the origin. The implication is that the distribution of inefficiency among firms can move away from zero. The gamma model is estimated by simulation methods either Bayesian MCMC [Huang (2003) and Tsionas (2002)] or maximum simulated likelihood [Greene (2003)]. Many other functional forms have been proposed. [See Greene (2008) for a survey.]

There are usually elements in the environment in which the firm operates that impact the firm's output and/or costs, but are not, themselves, outputs, inputs, or input prices. In our example 19.3, the costs of the Swiss railroads are affected by three variables, track width, long tunnels and curvature. It is not yet specified how such factors should be incorporated into the model; four candidates are in the mean and variance of  $u_i$ , directly in the function, or in the variance of  $v_i$ . [See Hadri, Guermat and Whittaker (2003) and Kumbhakar (1997c).] All of these can be found in the received studies. This aspect of the model was prominent in the discussion of the famous World Health Organization efficiency study of world health systems [WHO (2000), Tandon, Murray, Lauer and Evans (2000), and Greene (2004)]. In Example 19.3, we have placed the environmental factors in the mean of the inefficiency distribution. This produces a rather extreme set of results for the JLMS estimates of inefficiency - many railroads are estimated to be extremely inefficient. An alternative formulation would be a "heteroscedastic" model in which  $\sigma_{u,i} = \sigma_u \exp(z_i'\delta)$  or  $\sigma_{v,i} = \sigma_v \exp(z_i'\eta)$ , or both. We can see from the JLMS formula that the term heteroscedastic is actually a bit misleading, since both standard deviations enter (now)  $\lambda_i$ , which is, in turn, a crucial parameter in the mean of inefficiency.

How should inefficiency be modeled in panel data, such as in our example? It might be tempting to treat it as a time invariant "effect" (as in Schmidt and Sickles (1984) and Pitt and Lee (1984) in two pioneering papers). Greene (2004) argued that a preferable approach would be to allow inefficiency to vary freely over time in a panel, and to the extent that there is a common time invariant effect in the model, that should be treated as unobserved heterogeneity, not inefficiency. A string of studies, including Battese and Coelli (1992, 1995), Cuesta (2000),

Kumbhakar (1997a), Kumbhakar and Orea (2004), and many others have proposed hybrid forms that treat the core random part of inefficiency as a time invariant firm specific effect that is modified over time by a deterministic, possibly firm specific function. The Battese-Coelli form,

$$u_{it} = \exp[-\eta(t-T)]|U_i| \text{ where } U_i \sim N[0, \sigma_u^2],$$

has been used in a number of applications. Cuesta (2000) suggests allowing  $\eta$  to vary across firms, producing a model that bears some relationship to a fixed effects specification. This thread of the literature is one of the most active ongoing pursuits.

Is it reasonable to use a possibly restrictive parametric approach to modeling inefficiency? Sickles (2005) and Kumbhakar, Simar, Park and Tsionas (2007) are among numerous studies that have explored less parametric approaches to efficiency analysis. Proponents of "data envelopment analysis" [see, e.g., Simar and Wilson (2000, 2007)] have developed methods that impose absolutely no parametric structure on the production function. Among the costs of this high degree of flexibility is a difficulty to include environmental effects anywhere in the analysis, and the uncomfortable implication that any unmeasured heterogeneity of any sort is necessarily included in the measure of inefficiency. That is, data envelopment analysis returns to the deterministic frontier approach where this section began.

### Example 19.3 Stochastic Cost Frontier for Swiss Railroads

Farsi, Filippini and Greene (2005) analyzed the cost efficiency of Swiss railroads. In order to use the stochastic frontier approach to analyze costs of production, rather than production, we rely on the fundamental duality of production and cost [see Samuelson (1938), Shephard (1953), and Kumbhakar and Lovell (2000)]. An appropriate cost frontier model for a firm that produces more than one output — the Swiss railroads carry both freight and passengers — will appear as the following:

$$\ln(C/P_K) = \alpha + \sum_{k=1}^{K-1} \beta_k \ln(P_k/P_K) + \sum_{m=1}^M \gamma_m \ln Q_m + v + u.$$

The requirement that the cost function be homogeneous of degree one in the input prices has been imposed by normalizing total cost,  $C$ , and the first  $K-1$  prices by the  $K$ th input price. In this application, the three factors are labor, capital, and electricity — the third is used as the numeraire in the cost function. Notice that the inefficiency term,  $u$ , enters the cost function positively; actual cost is above the frontier cost. [The MLE is modified simply by replacing  $\varepsilon_i$  with  $-\varepsilon_i$  in (19-11).] In analyzing costs of production, we recognize that there is an additional source of inefficiency that is absent when we analyze production. On the production side, inefficiency measures the difference between output and frontier output, which arises because of technical inefficiency. By construction, if output fails to reach the efficient level for the given input usage, then costs must be higher than frontier costs. However, costs can be excessive even if the firm is technically efficient, if it is "allocatively inefficient." That is, the firm can be technically efficient while not using inputs in the cost minimizing mix (equating the ratio of marginal products to the input price ratios). It follows that on the cost side, " $u$ " can contain both elements of inefficiency while on the production side, we would expect to measure only technical inefficiency. [See Kumbhakar (1997b).]

thirty-seven

The data for this study are an unbalanced panel of 50 railroads with  $T_i$  ranging from 1 to 13. (37 of the firms are observed 13 times, 8 are observed 12 times, and the remaining 5 are observed 10, 7, 7, 3, and 1 times.) The variables we will use here are



- CT: Total costs adjusted for inflation (1000 Swiss franc),
- QP: Total passenger-output in passenger-kilometers,
- QF: Total goods-output in ton-kilometers,
- PL: Labor price adjusted for inflation (in Swiss Francs per person per year),
- PK: Capital price uses total number of seats as a proxy for capital stock (Swiss franc per seat)
- PE: Price of electricity (Swiss franc per kWh)

AV: Can you reword entry so it's not a sentence?

Logs of costs and prices ( $\ln CT$ ,  $\ln PK$ ,  $\ln PL$ ) are normalized by PE. We will also use these environmental variables:

- NARROW\_T: Dummy for the networks with narrow track (1 m wide) The usual width is 1.435m.
- TUNNEL: Dummy for networks that have tunnels with an average length of more than 300 meters.
- VIRAGE: Dummy for the networks whose minimum radius of curvature is 100 meters or less.

The full data set is given in Appendix Table F19.1. Several other variables not used here are presented in the Appendix Table. In what follows, we will ignore the panel data aspect of the data set. This would be a focal point of a more extensive study.

TB 19.1

There have been dozens of models proposed for the inefficiency component of the stochastic frontier model. Table 19.1 presents several different forms. The basic half-normal model is given in the first column. The estimated cost function parameters across the different forms are broadly similar, as might be expected as  $(\alpha, \beta)$  are consistently estimated in all cases. There are fairly pronounced differences in the implications for the components of  $\varepsilon$ , however.

There is an ambiguity in the model as to whether modifications to the distribution of  $u_i$  will affect the mean of the distribution, the variance, or both. The results below suggest that it is both for these data. The gamma and exponential models appear to remove most of the inefficiency from the data. Note that the estimates of  $\sigma_u$  are considerably smaller under these specifications, and  $\sigma_v$  is correspondingly larger. The second to last row shows the sample averages of the Jondrow estimators - this estimates  $E_\varepsilon E[u|\varepsilon] = E[u]$ . There is substantial difference across the specifications.

following/

The estimates in the rightmost two columns illustrate two different placements of the measured heterogeneity in the variance of  $u_i$  and directly in the cost function. The log likelihood function appears to favor the first of these. However, the models are not nested, and involve the same number of parameters. We used the Vuong test (see Section 14.6.6), instead, and obtained a value of -2.65 in favor of the heteroscedasticity model. Figure 19.3 describes the values of  $E[u|\varepsilon_i]$  estimated for the sample observations for the half-normal, heteroscedastic and heterogeneous models. The smaller estimate of  $\sigma_u$  for the third of these is evident in the figure, which suggests a somewhat tighter concentration of values than the other two.

minus

FIG 19.3

Table 19.1 Estimated Stochastic Frontier Cost Functions<sup>a</sup>

Variable	Model					
	Half Normal	Truncated Normal	Exponential	Gamma	Heterosced.	Heterogen.
Constant	-10.0799	-9.80624	-10.1838	-10.1944	-9.82189	-10.2891
lnQP	0.64220	0.62573	0.64403	0.64401	0.61976	0.63576
lnQF	0.06904	0.07708	0.06803	0.06810	0.07970	0.07526
lnPK	0.26005	0.26625	0.25883	0.25886	0.25464	0.25893
lnPL	0.53845	0.50474	0.56138	0.56047	0.53953	0.56036
Constant		0.44116			-2.48218 <sup>b</sup>	
Narrow		0.29881			2.16264 <sup>b</sup>	0.14355
Virage		-0.20738			-1.52964 <sup>b</sup>	-0.10483
Tunnel		0.01118			0.35748 <sup>b</sup>	-0.01914
$\sigma$	0.44240	0.38547	(0.34325)	(0.34288)	0.45392 <sup>c</sup>	0.40597
$\lambda$	1.27944	2.35055				0.91763
$P$			1.0000	1.22920		
$\theta$			13.2922	12.6915		
$\sigma_u$	(0.34857)	(0.35471)	(0.07523)	(0.09685)	0.37480 <sup>c</sup>	0.27448
$\sigma_v$	(0.27244)	(0.15090)	0.33490	0.33197	0.25606	0.29912
Mean $E[u \epsilon]$	0.27908	0.52858	0.075232	0.096616	0.29499	0.21926
lnL	-210.495	-200.67	-211.42	-211.091	-201.731	-208.349

<sup>a</sup> Estimates in parentheses are derived from other MLEs.  
<sup>b</sup> Estimates used in computation of  $\sigma_u$ .  
<sup>c</sup> Obtained by averaging  $\lambda = \sigma_u/\sigma_v$  over observations.

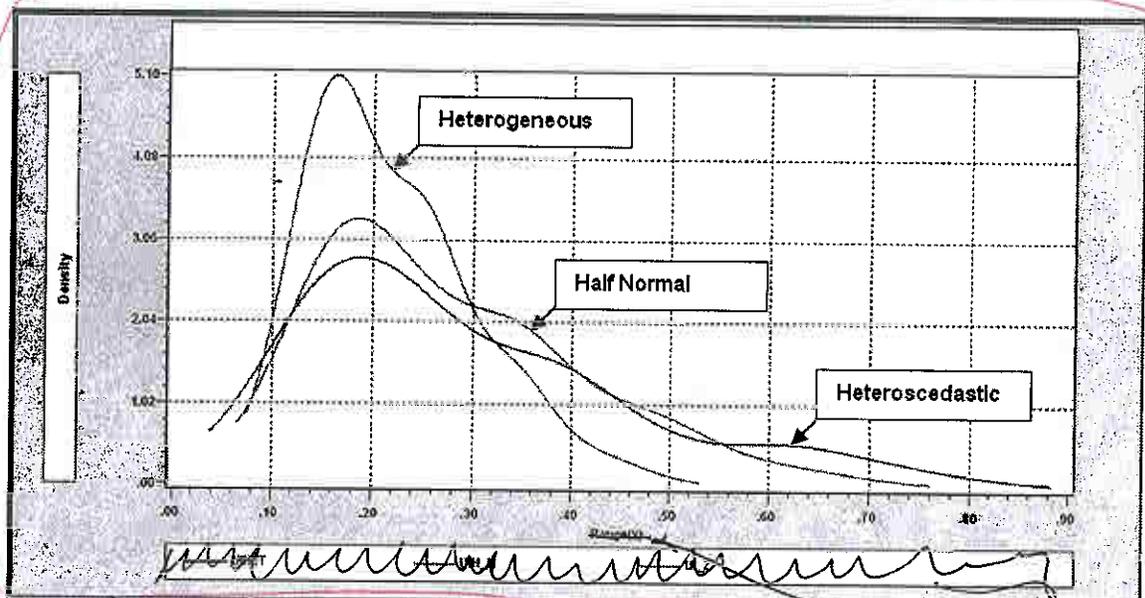


Figure 19.3 Kernel Density Estimator for JLMS Estimates

9

Inefficiency

19-16

CHAPTER 24 • Truncation, Censoring, and Sample Selection 869

of this model and maximum likelihood estimates.<sup>8</sup> The proportionality result appears to be quite general. In applications, it is usually found that, compared with consistent maximum likelihood estimates, the OLS estimates are biased toward zero. (See Example 24.4.)



19.3 ~~24.3~~ CENSORED DATA

A very common problem in microeconomic data is **censoring** of the dependent variable. When the dependent variable is censored, values in a certain range are all transformed to (or reported as) a single value. Some examples that have appeared in the empirical literature are as follows:

FN 9

1. Household purchases of durable goods [Tobin (1958)].
2. The number of extramarital affairs [Fair (1977, 1978)].
3. The number of hours worked by a woman in the labor force [Quester and Greene (1982)].
4. The number of arrests after release from prison [Witte (1980)].
5. Household expenditure on various commodity groups [Jarque (1987)].
6. Vacation expenditures [Melenberg and van Soest (1996)].

Av. Term "Censoring" was a RT on msp 19-1. Here also?

Each of these studies analyzes a dependent variable that is zero for a significant fraction of the observations. Conventional regression methods fail to account for the qualitative difference between limit (zero) observations and nonlimit (continuous) observations.

19.3.1 ~~24.4~~ THE CENSORED NORMAL DISTRIBUTION

The relevant distribution theory for a **censored variable** is similar to that for a truncated one. Once again, we begin with the normal distribution, as much of the received work has been based on an assumption of normality. We also assume that the censoring point is zero, although this is only a convenient normalization. In a truncated distribution, only the part of distribution above  $y = 0$  is relevant to our computations. To make the distribution integrate to one, we scale it up by the probability that an observation in the untruncated population falls in the range that interests us. When data are censored, the distribution that applies to the sample data is a mixture of discrete and continuous distributions. Figure 24.2 illustrates the effects.

FIG 19.4

19.4

To analyze this distribution, we define a new random variable  $y$  transformed from the original one,  $y^*$ , by

$$y = 0 \quad \text{if } y^* \leq 0,$$

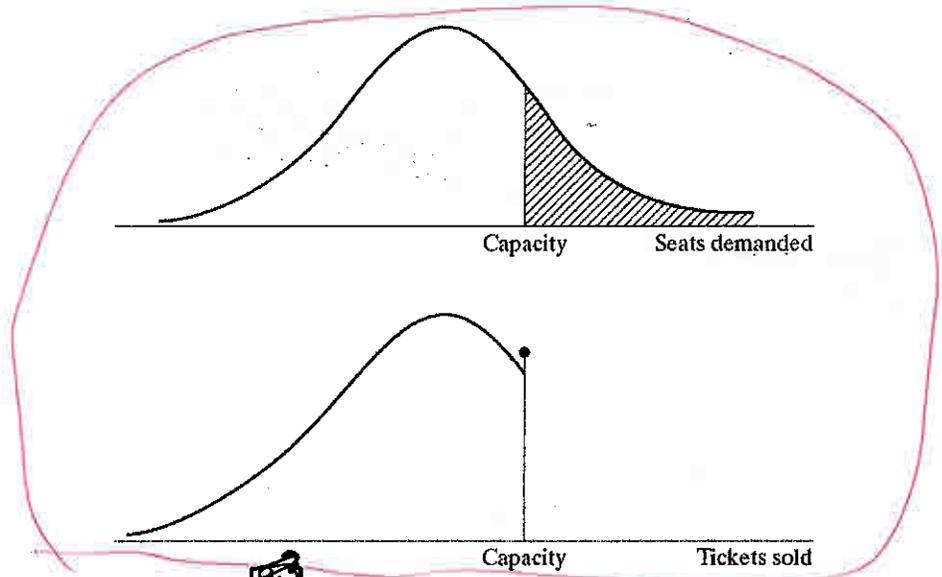
$$y = y^* \quad \text{if } y^* > 0.$$

The distribution that applies if  $y^* \sim N[\mu, \sigma^2]$  is  $\text{Prob}(y=0) = \text{Prob}(y^* \leq 0) = \Phi(-\mu/\sigma) = 1 - \Phi(\mu/\sigma)$ , and if  $y^* > 0$ , then  $y$  has the density of  $y^*$ .

This distribution is a mixture of discrete and continuous parts. The total probability is one, as required, but instead of scaling the second part, we simply assign the full probability in the censored region to the censoring point, in this case, zero.

See the appendix in Hausman and Wise (1977) and Greene (1985) as well. More extensive listings may be found in Amemiya (1984) and Maddala (1983).

870 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics



<sup>19.3</sup>  
~~FIGURE 24.2~~ Partially Censored Distribution.  
19.4

<sup>19.3</sup>  
~~THEOREM 24.3~~ Moments of the Censored Normal Variable

If  $y^* \sim N[\mu, \sigma^2]$  and  $y = a$  if  $y^* \leq a$  or else  $y = y^*$ , then

$$E[y] = \Phi a + (1 - \Phi)(\mu + \sigma\lambda),$$

and

$$\text{Var}[y] = \sigma^2(1 - \Phi)[(1 - \delta) + (\alpha - \lambda)^2\Phi],$$

where

$$\Phi[(a - \mu)/\sigma] = \Phi(\alpha) = \text{Prob}(y^* \leq a) = \Phi, \quad \lambda = \phi/(1 - \Phi),$$

and

$$\delta = \lambda^2 - \lambda\alpha.$$

**Proof:** For the mean,

$$\begin{aligned} E[y] &= \text{Prob}(y = a) \times E[y | y = a] + \text{Prob}(y > a) \times E[y | y > a] \\ &= \text{Prob}(y^* \leq a) \times a + \text{Prob}(y^* > a) \times E[y^* | y^* > a] \\ &= \Phi a + (1 - \Phi)(\mu + \sigma\lambda) \end{aligned}$$

19.2

using Theorem ~~24.2~~. For the variance, we use a counterpart to the decomposition in (B-69), that is,  $\text{Var}[y] = E[\text{conditional variance}] + \text{Var}[\text{conditional mean}]$ , and Theorem ~~24.2~~.

19.2

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 871

For the special case of  $a = 0$ , the mean simplifies to

$$E[y|a = 0] = \Phi(\mu/\sigma)(\mu + \sigma\lambda), \quad \text{where } \lambda = \frac{\phi(\mu/\sigma)}{\Phi(\mu/\sigma)}$$

For censoring of the upper part of the distribution instead of the lower, it is only necessary to reverse the role of  $\Phi$  and  $1 - \Phi$  and redefine  $\lambda$  as in Theorem ~~24.2~~.

<sup>19.4</sup> ~~24.3~~ **Example 19.4 Censored Random Variable**

We are interested in the number of tickets demanded for events at a certain arena. Our only measure is the number actually sold. Whenever an event sells out, however, we know that the actual number demanded is larger than the number sold. The number of tickets demanded is censored when it is transformed to obtain the number sold. Suppose that the arena in question has 20,000 seats and, in a recent season, sold out 25 percent of the time. If the average attendance, including sellouts, was 18,000, then what are the mean and standard deviation of the demand for seats? According to Theorem ~~24.3~~, the 18,000 is an estimate of

$$E[\text{sales}] = 20,000(1 - \Phi) + [\mu + \sigma\lambda]\Phi.$$

Because this is censoring from above, rather than below,  $\lambda = -\phi(\alpha)/\Phi(\alpha)$ . The argument of  $\Phi$ ,  $\phi$ , and  $\lambda$  is  $\alpha = (20,000 - \mu)/\sigma$ . If 25 percent of the events are sellouts, then  $\Phi = 0.75$ . Inverting the standard normal at 0.75 gives  $\alpha = 0.675$ . In addition, if  $\alpha = 0.675$ , then  $-\phi(0.675)/0.75 = \lambda = -0.424$ . This result provides two equations in  $\mu$  and  $\sigma$ , (a)  $18,000 = 0.25(20,000) + 0.75(\mu - 0.424\sigma)$  and (b)  $0.675\sigma = 20,000 - \mu$ . The solutions are  $\sigma = 2426$  and  $\mu = 18,362$ .

For comparison, suppose that we were told that the mean of 18,000 applies only to the events that were not sold out and that, on average, the arena sells out 25 percent of the time. Now our estimates would be obtained from the equations (a)  $18,000 = \mu - 0.424\sigma$  and (b)  $0.675\sigma = 20,000 - \mu$ . The solutions are  $\sigma = 1820$  and  $\mu = 18,772$ .

19.3.2 ~~24.3.2~~ **THE CENSORED REGRESSION (TOBIT) MODEL**

The regression model based on the preceding discussion is referred to as the censored regression model or the tobit model. [In reference to Tobin (1958), where the model was first proposed.] The regression is obtained by making the mean in the preceding correspond to a classical regression model. The general formulation is usually given in terms of an index function,

$$\begin{aligned} y_i^* &= x_i'\beta + \varepsilon_i, \\ y_i &= 0 \quad \text{if } y_i^* \leq 0, \\ y_i &= y_i^* \quad \text{if } y_i^* > 0. \end{aligned}$$

There are potentially three conditional mean functions to consider, depending on the purpose of the study. For the index variable, sometimes called the latent variable,  $E[y_i^* | x_i]$  is  $x_i'\beta$ . If the data are always censored, however, then this result will usually not be useful. Consistent with Theorem ~~24.3~~, for an observation randomly drawn from the population, which may or may not be censored,

$$E[y_i | x_i] = \Phi\left(\frac{x_i'\beta}{\sigma}\right)(x_i'\beta + \sigma\lambda_i),$$

where

$$\lambda_i = \frac{\phi[(0 - x_i'\beta)/\sigma]}{1 - \Phi[(0 - x_i'\beta)/\sigma]} = \frac{\phi(x_i'\beta/\sigma)}{\Phi(x_i'\beta/\sigma)}$$

19  
(24-12)

872 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

Finally, if we intend to confine our attention to uncensored observations, then the results for the truncated regression model apply. The limit observations should not be discarded, however, because the truncated regression model is no more amenable to least squares than the censored data model. It is an unresolved question which of these functions should be used for computing predicted values from this model. Intuition suggests that  $E[y_i | x_i]$  is correct, but authors differ on this point. For the setting in Example 24.3, for predicting the number of tickets sold, say, to plan for an upcoming event, the censored mean is obviously the relevant quantity. On the other hand, if the objective is to study the need for a new facility, then the mean of the latent variable  $y_i^*$  would be more interesting.

19.4

There are differences in the <sup>partial</sup> ~~marginal~~ effects as well. For the index variable.

$$\frac{\partial E[y_i^* | x_i]}{\partial x_i} = \beta.$$

But this result is not what will usually be of interest, because  $y_i^*$  is unobserved. For the observed data,  $y_i$ , the following general result will be useful.

FN 10

**19.4 Partial**  
**THEOREM 24.4 Marginal Effects in the Censored Regression Model**

In the censored regression model with latent regression  $y^* = x'\beta + \varepsilon$  and observed dependent variable,  $y = a$  if  $y^* \leq a$ ,  $y = b$  if  $y^* \geq b$ , and  $y = y^*$  otherwise, where  $a$  and  $b$  are constants, let  $f(\varepsilon)$  and  $F(\varepsilon)$  denote the density and cdf of  $\varepsilon$ . Assume that  $\varepsilon$  is a continuous random variable with mean 0 and variance  $\sigma^2$ , and  $f(\varepsilon | x) = f(\varepsilon)$ . Then

$$\frac{\partial E[y | x]}{\partial x} = \beta \times \text{Prob}[a < y^* < b].$$

*Proof:* By definition,

$$E[y | x] = a \text{Prob}[y^* \leq a | x] + b \text{Prob}[y^* \geq b | x] + \text{Prob}[a < y^* < b | x] E[y^* | a < y^* < b | x].$$

Let  $\alpha_j = (j - x'\beta)/\sigma$ ,  $F_j = F(\alpha_j)$ ,  $f_j = f(\alpha_j)$ , and  $j = a, b$ . Then

$$E[y | x] = aF_a + b(1 - F_b) + (F_b - F_a) E[y^* | a < y^* < b, x].$$

Because  $y^* = x'\beta + \sigma[(y^* - \beta'x)/\sigma]$ , the conditional mean may be written

$$E[y^* | a < y^* < b, x] = x'\beta + \sigma E\left[\frac{y^* - x'\beta}{\sigma} \mid \frac{a - x'\beta}{\sigma} < \frac{y^* - x'\beta}{\sigma} < \frac{b - x'\beta}{\sigma}\right] \\ = x'\beta + \sigma \int_{\alpha_a}^{\alpha_b} \frac{(\varepsilon/\sigma) f(\varepsilon/\sigma)}{F_b - F_a} d\left(\frac{\varepsilon}{\sigma}\right).$$

See Greene (1999) for the general result and Rosett and Nelson (1975) and Nakamura and Nakamura (1983) for applications based on the normal distribution.

874 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

19.2

TABLE 24.71 Tobit Estimates of an Hours Worked Equation

	White Wives		Black Wives		Least Squares	Scaled OLS
	Coefficient	Slope	Coefficient	Slope		
Constant	-1803.13 (-8.64)		-2753.87 (-9.68)			
Small kids	-1324.84 (-19.78)	-385.89	-824.19 (-10.14)	-376.53	-352.63	-766.56
Education difference	-48.08 (-4.77)	-14.00	22.59 (1.96)	10.32	11.47	24.93
Relative wage	312.07 (5.71)	90.90	286.39 (3.32)	130.93	123.95	269.46
Second marriage	175.85 (3.47)	51.51	25.33 (0.41)	11.57	13.14	28.57
Mean divorce probability	417.39 (6.52)	121.58	481.02 (5.28)	219.75	219.22	476.57
High divorce probability	670.22 (8.40)	195.22	578.66 (5.33)	264.36	244.17	530.80
$\sigma$	1559	618	1511	826		
Sample size	7459		2798			
Proportion working	0.29		0.46			

table implies that a very large proportion of the women reported zero hours, so least squares regression would be inappropriate.

The figures in parentheses are the ratio of the coefficient estimate to the estimated asymptotic standard error. The dependent variable is hours worked in the survey year. "Small kids" is a dummy variable indicating whether there were children in the household. The "education difference" and "relative wage" variables compare husband and wife on these two dimensions. The wage rate used for wives was predicted using a previously estimated regression model and is thus available for all individuals, whether working or not. "Second marriage" is a dummy variable. Divorce probabilities were produced by a large microsimulation model presented in another study [Orcutt, Caldwell, and Wertheimer (1976)]. The variables used here were dummy variables indicating "mean" if the predicted probability was between 0.01 and 0.03 and "high" if it was greater than 0.03. The "slopes" are the marginal effects described earlier.

Note the marginal effects compared with the tobit coefficients. Likewise, the estimate of  $\sigma$  is quite misleading as an estimate of the standard deviation of hours worked.

The effects of the divorce probability variables were as expected and were quite large. One of the questions raised in connection with this study was whether the divorce probabilities could reasonably be treated as independent variables. It might be that for these individuals, the number of hours worked was a significant determinant of the probability.

19.3.3 ESTIMATION

Estimation of this model is very similar to that of truncated regression. The tobit model has become so routine and been incorporated in so many computer packages that despite formidable obstacles in years past, estimation is now essentially on the level of ordinary linear regression. The log-likelihood for the censored regression model is

$$\ln L = \sum_{y_i > 0} -\frac{1}{2} \left[ \log(2\pi) + \ln \sigma^2 + \frac{(y_i - x_i' \beta)^2}{\sigma^2} \right] + \sum_{y_i = 0} \ln \left[ 1 - \Phi \left( \frac{x_i' \beta}{\sigma} \right) \right] \quad (24.13)$$

*Handwritten scribble*

**19.4**  
**THEOREM 24.4** (Continued)  
 Collecting terms, we have

$$E[y|x] = aF_a + b(1 - F_b) + (F_b - F_a)\beta'x + \sigma \int_{\alpha_a}^{\alpha_b} \left(\frac{\varepsilon}{\sigma}\right) f\left(\frac{\varepsilon}{\sigma}\right) d\left(\frac{\varepsilon}{\sigma}\right).$$

Now, differentiate with respect to  $x$ . The only complication is the last term, for which the differentiation is with respect to the limits of integration. We use Leibnitz's theorem and use the assumption that  $f(\varepsilon)$  does not involve  $x$ . Thus,

$$\frac{\partial E[y|x]}{\partial x} = \left(\frac{-\beta}{\sigma}\right)af_a - \left(\frac{-\beta}{\sigma}\right)bf_b + (F_b - F_a)\beta + (x'\beta)(f_b - f_a)\left(\frac{-\beta}{\sigma}\right) + \sigma[\alpha_b f_b - \alpha_a f_a]\left(\frac{-\beta}{\sigma}\right).$$

After inserting the definitions of  $\alpha_a$  and  $\alpha_b$ , and collecting terms, we find all terms sum to zero save for the desired result,

$$\frac{\partial E[y|x]}{\partial x} = (F_b - F_a)\beta = \beta \times \text{Prob}[a < y_i^* < b].$$

Note that this general result includes censoring in either or both tails of the distribution, and it does not assume that  $\varepsilon$  is normally distributed. For the standard case with censoring at zero and normally distributed disturbances, the result specializes to

$$\frac{\partial E[y_i|x_i]}{\partial x_i} = \beta\Phi\left(\frac{x_i'\beta}{\sigma}\right).$$

Although not a formal result, this does suggest a reason why, in general, least squares estimates of the coefficients in a tobit model usually resemble the MLEs times the proportion of nonlimit observations in the sample.

McDonald and Moffitt (1980) suggested a useful decomposition of  $\partial E[y_i|x_i]/\partial x_i$ ,

$$\frac{\partial E[y_i|x_i]}{\partial x_i} = \beta \times \{ \Phi_i [1 - \lambda_i(\alpha_i + \lambda_i)] + \phi_i(\alpha_i + \lambda_i) \},$$

where  $\alpha_i = x_i'\beta/\sigma$ ,  $\Phi_i = \Phi(\alpha_i)$  and  $\lambda_i = \phi_i/\Phi_i$ . Taking the two parts separately, this result decomposes the slope vector into

$$\frac{\partial E[y_i|x_i]}{\partial x_i} = \text{Prob}[y_i > 0] \frac{\partial E[y_i|x_i, y_i > 0]}{\partial x_i} + E[y_i|x_i, y_i > 0] \frac{\partial \text{Prob}[y_i > 0]}{\partial x_i}.$$

Thus, a change in  $x_i$  has two effects: It affects the conditional mean of  $y_i^*$  in the positive part of the distribution, and it affects the probability that the observation will fall in that part of the distribution.

**19.5**  
**Example 24.4** *Estimated Tobit Equations for Hours Worked*  
 In their study of the number of hours worked in a survey year by a large sample of wives, Quester and Greene (1982) were interested in whether wives whose marriages were statistically more likely to dissolve hedged against that possibility by spending, on average, more time working. They reported the tobit estimates given in Table 24.7. The last figure in the

19.2

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 875

The two parts correspond to the classical regression for the nonlimit observations and the relevant probabilities for the limit observations, respectively. This likelihood is a nonstandard type, because it is a mixture of discrete and continuous distributions. In a seminal paper, Amemiya (1973) showed that despite the complications, proceeding in the usual fashion to maximize  $\ln L$  would produce an estimator with all the familiar desirable properties attained by MLEs.

The log-likelihood function is fairly involved, but **Olsen's (1978) reparameterization** simplifies things considerably. With  $\gamma = \beta/\sigma$  and  $\theta = 1/\sigma$ , the log-likelihood is

$$\ln L = \sum_{y_i > 0} -\frac{1}{2} [\ln(2\pi) - \ln \theta^2 + (\theta y_i - x_i' \gamma)^2] + \sum_{y_i = 0} \ln[1 - \Phi(x_i' \gamma)] \quad (19-14)$$

The results in this setting are now very similar to those for the truncated regression. The Hessian is always negative definite, so Newton's method is simple to use and usually converges quickly. After convergence, the original parameters can be recovered using  $\sigma = 1/\theta$  and  $\beta = \gamma/\theta$ . The asymptotic covariance matrix for these estimates can be obtained from that for the estimates of  $[\gamma, \theta]$  using the **delta method**; Est. Asy. Var $[\hat{\beta}, \hat{\sigma}] = \mathbf{J} \text{Asy. Var}[\hat{\gamma}, \hat{\theta}] \mathbf{J}'$ , where

$$\mathbf{J} = \begin{bmatrix} \partial \beta / \partial \gamma' & \partial \beta / \partial \theta \\ \partial \sigma / \partial \gamma' & \partial \sigma / \partial \theta \end{bmatrix} = \begin{bmatrix} (1/\theta) \mathbf{I} & (-1/\theta^2) \gamma \\ \mathbf{0} & (-1/\theta^2) \end{bmatrix}$$

Researchers often compute ordinary least squares estimates despite their inconsistency. Almost without exception, it is found that the OLS estimates are smaller in absolute value than the MLEs. A striking empirical regularity is that the maximum likelihood estimates can often be approximated by dividing the OLS estimates by the proportion of nonlimit observations in the sample.<sup>12</sup> The effect is illustrated in the last two columns of Table 24.1. Another strategy is to discard the limit observations, but we now see that just trades the censoring problem for the truncation problem.

FN 11  
TB 19.2

13

19.3.4 ~~24.3.4~~ SOME ISSUES IN SPECIFICATION

Two issues that commonly arise in microeconomic data, heteroscedasticity and nonnormality, have been analyzed at length in the tobit setting.<sup>13</sup>

19.3.4.a ~~24.3.4.a~~ Heteroscedasticity

Maddala and Nelson (1975), Hurd (1979), Arabmazar and Schmidt (1982a,b), and Brown and Moffitt (1982) all have varying degrees of pessimism regarding how inconsistent the maximum likelihood estimator will be when heteroscedasticity occurs. Not surprisingly, the degree of censoring is the primary determinant. Unfortunately, all the analyses have been carried out in the setting of very specific models—for example, involving only a single dummy variable or one with groupwise heteroscedasticity—so the primary lesson is the very general conclusion that heteroscedasticity emerges as an obviously serious problem.

<sup>13</sup> This concept is explored further in Greene (1980b), Goldberger (1981), and Chung and Goldberger (1984).  
<sup>14</sup> Two symposia that contain numerous results on these subjects are Blundell (1987) and Durkin (1986b). An application that explores these two issues in detail is Melchberg and van Soest (1996).

**19.4**  
**THEOREM 24.4 (Continued)**

Collecting terms, we have

$$E[y|\mathbf{x}] = aF_a + b(1 - F_b) + (F_b - F_a)\beta'x + \sigma \int_{a_a}^{\alpha_b} \left(\frac{\varepsilon}{\sigma}\right) f\left(\frac{\varepsilon}{\sigma}\right) d\left(\frac{\varepsilon}{\sigma}\right).$$

Now, differentiate with respect to  $\mathbf{x}$ . The only complication is the last term, for which the differentiation is with respect to the limits of integration. We use Leibnitz's theorem and use the assumption that  $f(\varepsilon)$  does not involve  $\mathbf{x}$ . Thus,

$$\begin{aligned} \frac{\partial E[y|\mathbf{x}]}{\partial \mathbf{x}} &= \left(\frac{-\beta'}{\sigma}\right) a f_a - \left(\frac{-\beta'}{\sigma}\right) b f_b + (F_b - F_a)\beta + (\mathbf{x}'\beta)(f_b - f_a) \left(\frac{-\beta}{\sigma}\right) \\ &\quad + \sigma[\alpha_b f_b - \alpha_a f_a] \left(\frac{-\beta}{\sigma}\right). \end{aligned}$$

After inserting the definitions of  $\alpha_a$  and  $\alpha_b$ , and collecting terms, we find all terms sum to zero save for the desired result,

$$\frac{\partial E[y|\mathbf{x}]}{\partial \mathbf{x}} = (F_b - F_a)\beta = \beta \times \text{Prob}[a < y_i^* < b].$$

Note that this general result includes censoring in either or both tails of the distribution, and it does not assume that  $\varepsilon$  is normally distributed. For the standard case with censoring at zero and normally distributed disturbances, the result specializes to

$$\frac{\partial E[y_i|\mathbf{x}_i]}{\partial \mathbf{x}_i} = \beta \Phi\left(\frac{\mathbf{x}_i'\beta}{\sigma}\right).$$

Although not a formal result, this does suggest a reason why, in general, least squares estimates of the coefficients in a tobit model usually resemble the MLEs times the proportion of nonlimit observations in the sample.

McDonald and Moffitt (1980) suggested a useful decomposition of  $\partial E[y_i|\mathbf{x}_i]/\partial \mathbf{x}_i$ ,

$$\frac{\partial E[y_i|\mathbf{x}_i]}{\partial \mathbf{x}_i} = \beta \times \{ \Phi_i [1 - \lambda_i(\alpha_i + \lambda_i)] + \phi_i(\alpha_i + \lambda_i) \},$$

where  $\alpha_i = \mathbf{x}_i'\beta/\sigma$ ,  $\Phi_i = \Phi(\alpha_i)$  and  $\lambda_i = \phi_i/\Phi_i$ . Taking the two parts separately, this result decomposes the slope vector into

$$\frac{\partial E[y_i|\mathbf{x}_i]}{\partial \mathbf{x}_i} = \text{Prob}[y_i > 0] \frac{\partial E[y_i|\mathbf{x}_i, y_i > 0]}{\partial \mathbf{x}_i} + E[y_i|\mathbf{x}_i, y_i > 0] \frac{\partial \text{Prob}[y_i > 0]}{\partial \mathbf{x}_i}.$$

Thus, a change in  $\mathbf{x}_i$  has two effects: It affects the conditional mean of  $y_i^*$  in the positive part of the distribution, and it affects the probability that the observation will fall in that part of the distribution.

**19.5**  
**Example 24.4 Estimated Tobit Equations for Hours Worked**

In their study of the number of hours worked in a survey year by a large sample of wives, Quester and Greene (1982) were interested in whether wives whose marriages were statistically more likely to dissolve hedged against that possibility by spending, on average, more time working. They reported the tobit estimates given in Table 24.4. The last figure in the

19-24

874 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

19.2

TABLE 24.7 Tobit Estimates of an Hours Worked Equation

	White Wives		Black Wives		Least Squares	Scaled OLS
	Coefficient	Slope	Coefficient	Slope		
Constant	-1803.13 (-8.64)		-2753.87 (-9.68)			
Small kids	-1324.84 (-19.78)	-385.89	-824.19 (-10.14)	-376.53	-352.63	-766.56
Education difference	-48.08 (-4.77)	-14.00	22.59 (1.96)	10.32	11.47	24.93
Relative wage	312.07 (5.71)	90.90	286.39 (3.32)	130.93	123.95	269.46
Second marriage	175.85 (3.47)	51.51	25.33 (0.41)	11.57	13.14	28.57
Mean divorce probability	417.39 (6.52)	121.58	481.02 (5.28)	219.75	219.22	476.57
High divorce probability	670.22 (8.40)	195.22	578.66 (5.33)	264.36	244.17	530.80
$\sigma$	1559	618	1511	826		
Sample size		7459		2798		
Proportion working		0.29		0.46		

table implies that a very large proportion of the women reported zero hours, so least squares regression would be inappropriate.

The figures in parentheses are the ratio of the coefficient estimate to the estimated asymptotic standard error. The dependent variable is hours worked in the survey year. "Small kids" is a dummy variable indicating whether there were children in the household. The "education difference" and "relative wage" variables compare husband and wife on these two dimensions. The wage rate used for wives was predicted using a previously estimated regression model and is thus available for all individuals, whether working or not. "Second marriage" is a dummy variable. Divorce probabilities were produced by a large microsimulation model presented in another study [Orcutt, Caldwell, and Wertheimer (1976)]. The variables used here were dummy variables indicating "mean" if the predicted probability was between 0.01 and 0.03 and "high" if it was greater than 0.03. The "slopes" are the marginal effects described earlier.

Note the marginal effects compared with the tobit coefficients. Likewise, the estimate of  $\sigma$  is quite misleading as an estimate of the standard deviation of hours worked.

The effects of the divorce probability variables were as expected and were quite large. One of the questions raised in connection with this study was whether the divorce probabilities could reasonably be treated as independent variables. It might be that for these individuals, the number of hours worked was a significant determinant of the probability.

19.3.3 ESTIMATION

~~Estimation of this model is very similar to that of truncated regression.~~ The tobit model has become so routine and been incorporated in so many computer packages that despite formidable obstacles in years past, estimation is now essentially on the level of ordinary linear regression. The log-likelihood for the censored regression model is

$$\ln L = \sum_{y_i > 0} -\frac{1}{2} \left[ \log(2\pi) + \ln \sigma^2 + \frac{(y_i - x_i' \beta)^2}{\sigma^2} \right] + \sum_{y_i = 0} \ln \left[ 1 - \Phi \left( \frac{x_i' \beta}{\sigma} \right) \right] \quad 19 \quad (24.13)$$

~~Not a tobit~~

### 19.3.4 Two Part Models and Corner Solutions

The Tobit Model contains a restriction that might be unreasonable in an economic setting. Consider a behavioral outcome,  $y$  = charitable donation. Two implications of the tobit model are that

$$\text{Prob}(y > 0 | \mathbf{x}) = \text{Prob}(\mathbf{x}'\boldsymbol{\beta} + \varepsilon > 0 | \mathbf{x}) = \Phi(\mathbf{x}'\boldsymbol{\beta}/\sigma)$$

and [from (19-7)]

$$E[y | y > 0, \mathbf{x}] = \mathbf{x}'\boldsymbol{\beta} + \sigma\phi(\mathbf{x}'\boldsymbol{\beta}/\sigma)/\Phi(\mathbf{x}'\boldsymbol{\beta}/\sigma).$$

Differentiating both of these, we find from (17-11) and (19-8),

$$\partial \text{Prob}(y > 0 | \mathbf{x}) / \partial \mathbf{x} = [\phi(\mathbf{x}'\boldsymbol{\beta}/\sigma)/\sigma]\boldsymbol{\beta} = \text{a positive multiple of } \boldsymbol{\beta},$$

$$\partial E[y | y > 0, \mathbf{x}] / \partial \mathbf{x} = \{[1 - \delta(\mathbf{x}'\boldsymbol{\beta}/\sigma)]/\sigma\}\boldsymbol{\beta} = \text{a positive multiple of } \boldsymbol{\beta}.$$

Thus, any variable that appears in the model affects the participation probability and the intensity equation with the same sign. In the case suggested, for example, it is conceivable that age might affect participation and intensity in different directions. Fin and Schmidt (1984) suggest another application, loss due to fire in buildings; older buildings might be more likely to have fires but, because of the greater value of newer buildings, the actual damage might be greater in newer buildings. This fact would require the coefficient on age to have different signs in the two functions, which is impossible in the tobit model because they are the same coefficient.

In an early study in this literature, Cragg (1971) proposed a somewhat more general model in which the probability of a limit observation is independent of the regression model for the nonlimit data. One can imagine, for instance, the decision of whether or not to purchase a car as being different from the decision of how much to spend on the car, having decided to buy one. A more general model that accommodates these objections is as follows:

#### 1. Participation equation:

$$\text{Prob}[y_i^* > 0] = \Phi(\mathbf{x}_i'\boldsymbol{\gamma}), \quad d_i = 1 \text{ if } y_i^* > 0,$$

$$\text{Prob}[y_i^* \leq 0] = 1 - \Phi(\mathbf{x}_i'\boldsymbol{\gamma}), \quad d_i = 0 \text{ if } y_i^* \leq 0.$$

(19-15)

#### 2. Intensity equation for nonlimit observations:

$$E[y_i | d_i = 1] = \mathbf{x}_i'\boldsymbol{\beta} + \sigma\lambda_i,$$

according to Theorem 19.2. This two model is a combination of the truncated regression model of Section 24.2 and the univariate probit model of Section 17.3, which suggests a method of analyzing it. Note that it is precisely the same approach we considered in Section 18.4.8 and Example 18.12 where we used a hurdle model to model doctor visits. The tobit model returns if  $y = \boldsymbol{\beta}/\sigma$ . The parameters of the regression (intensity) equation can be estimated independently using the truncated regression model of Section 19.2. An application is Melenberg and van Soest (1996).

Lin and Schmidt (1984) considered testing the restriction of the tobit model. Based only on the tobit model, they devised a Lagrange multiplier statistic that, although a bit cumbersome algebraically, can be computed without great difficulty. If one is able to estimate the truncated regression model, the tobit model, and the probit model separately, then there is a simpler way to test the hypothesis. The tobit log-likelihood is the sum of the log-likelihoods for the truncated

AO: Confirm  
x-ref to  
Sec 24.2  
is OK

Ami confirm  
x-refs to  
(24-13) &  
(23-20)  
are OK

regression and probit models. To show this result, add and subtract  $\sum_{y_i=1} \ln \Phi(\mathbf{x}_i' \boldsymbol{\beta})$  in (24-13). This produces the loglikelihood for the truncated regression model (considered in the exercises) plus (23-20) for the probit model. Therefore, a likelihood ratio statistic can be computed using

$$\lambda = -2[\ln LT - (\ln LP + \ln LTR)],$$

where

- $LT$  = likelihood for the tobit model in (19-13), with the same coefficients,
- $LP$  = likelihood for the probit model in (17-17), fit separately,
- $LTR$  = likelihood for the truncated regression model, fit separately.

preceding/

The two part model just considered extends the tobit model, but it stops a bit short of the generality we might achieve. In the hurdle model above, we have assumed that the same regressors appear in both equations. Although this produces a convenient way to retreat to the tobit model as a parametric restriction, it couples the two decisions perhaps unreasonably. In our example to follow, where we model extramarital affairs, the decision whether or not to spend any time in an affair may well be an entirely different decision from how much time to spend having once made that commitment. The obvious way to proceed is to reformulate the hurdle model as

1. Participation equation:

$$\begin{aligned} \text{Prob}[d_i^* > 0] &= \Phi(\mathbf{z}_i' \boldsymbol{\gamma}), & d_i &= 1 \text{ if } d_i^* > 0, \\ \text{Prob}[d_i^* \leq 0] &= 1 - \Phi(\mathbf{z}_i' \boldsymbol{\gamma}), & d_i &= 0 \text{ if } d_i^* \leq 0. \end{aligned} \quad (19-16)$$

2. Intensity equation for nonlimit observations:

$$E[y_i | d_i = 1] = \mathbf{x}_i' \boldsymbol{\beta} + \sigma \lambda_i.$$

This extension, however, omits an important element; it seems unlikely that the two decisions would be uncorrelated; that is, the implicit disturbances in the equations should be correlated. The combination of these produces what has been labeled a "Type II tobit model." [Amemiya (1985) identified five possible permutations of the model specification and observation mechanism. The familiar tobit model is Type I; this is Type II.] The full model is

KT

1. Participation equation

$$\begin{aligned} d_i^* &= \mathbf{z}_i' \boldsymbol{\gamma} + u_i, \quad u_i \sim N[0, 1] \\ d_i &= 1 \text{ if } d_i^* > 0, \quad 0 \text{ otherwise} \end{aligned}$$

2. Intensity equation

$$y_i^* = \mathbf{x}_i' \boldsymbol{\beta} + \varepsilon_i, \quad \varepsilon_i \sim N[0, \sigma^2]$$

3. Observation mechanism

- (a)  $y_i^* = 0$  if  $d_i = 0$  and  $y_i = y_i^*$  if  $d_i = 1$
- (b)  $y_i = y_i^*$  if  $d_i = 1$  and  $y_i$  is unobserved if  $d_i = 0$ .

4. Endogeneity

$$(u_i, \varepsilon_i) \sim \text{bivariate normal with correlation } \rho.$$

Mechanism (a) produces Amemiya's Type II model. (Amemiya blends these two interpretations. In the statement of the model, he presents (a), but in the subsequent discussion, assumes (b). The difference is substantive if  $\mathbf{x}_i$  is observed in case (b). Otherwise, they are the same, and " $y_i = 0$ " is not actually meaningful. Amemiya notes, " $y_i^* = 0$  merely signifies the event  $d_i^* \leq 0$ ." If  $\mathbf{x}_i$  is observed when  $d_i = 0$ , these observations will contribute to the likelihood for the full sample. If not, then they will not. We will develop this idea later when we consider Heckman's selection model (which is case (b) without observed  $\mathbf{x}_i$  when  $d_i = 0$ ).

then/

There are two estimation strategies that can be used to fit the Type II model. A two step method can proceed as follows: The probit model for  $d_i$  can be estimated using maximum likelihood as shown in Section 17.3. For the second step, we make use of our theorems on truncation (and Theorem 19.5 that will appear later) to write

$$\begin{aligned}
E[y_i | d_i = 1, x_i, z_i] &= x_i' \beta + E[\varepsilon_i | d_i = 1, x_i, z_i] \\
&= x_i' \beta + \rho \sigma \frac{\phi(z_i' \gamma)}{\Phi(z_i' \gamma)} \\
&= x_i' \beta + \rho \sigma \lambda_i
\end{aligned}
\tag{19-17}$$

Since we have estimated  $\gamma$  at step 1, we can compute  $\hat{\lambda}_i = \phi(z_i' \hat{\gamma}) / \Phi(z_i' \hat{\gamma})$  using the first step estimates, <sup>and</sup> we can estimate  $\beta$  and  $\theta = (\rho\sigma)$  by least squares regression of  $y_i$  on  $x_i$  and  $\hat{\lambda}_i$ . It will be necessary to correct the asymptotic covariance matrix that is computed for  $(\hat{\beta}, \hat{\theta})$ . This is a template application of the Murphy and Topel (2002) results that appear in Section 14.7. The second approach is full information maximum likelihood, estimating all of the parameters in both equations simultaneously. We will return to the details of estimation of the Type II tobit model in Section 19.5 where we examine Heckman's model of "sample selection" model (which is the Type II tobit model).

Many of the applications of the tobit model in the received literature are constructed not to accommodate censoring of the underlying data, but, rather, to model the appearance of a large cluster of zeros. Cragg's application is clearly related to this phenomenon. Consider, for example, survey data on purchases of consumer durables, firm expenditure on research and development, or consumer savings. In each case, the observed data will consist of zero or some positive amount. Arguably, there are two decisions at work in these scenarios: first, whether to engage in the activity or not, and second, given that the answer to the first question is yes, how intensively to engage in it—how much to spend, for example. This is precisely the motivation behind the hurdle model. This specification has been labeled a "**corner solution model**"; see Wooldridge (2002a, pp. 518-519).

In practical terms, the difference between the hurdle model and the tobit model should be evident in the data. Often overlooked in tobit analyses is that the model predicts not only a cluster of zeros (or limit observations), but also a grouping of observations *near zero* (or the limit point). For example, the tobit model is surely misspecified for the sort of (hypothetical) spending data shown in Figure 19.5 for a sample of 1,000 observations. Neglecting for the moment the earlier point about the underlying decision process, Figure 19.6 shows the characteristic appearance of a (substantively) censored variable. The implication for the model builder is that an appropriate specification would consist of two equations, one for the "participation decision," and one for the distribution of the positive dependent variable. Formally, we might, continuing the development of Cragg's specification, model the first decision with a binary choice (e.g., probit or logit model). The second equation is a model for  $y | y > 0$ , for which the truncated regression model of Section 19.2.3 is a natural candidate. As we will see, this is essentially the model behind the sample selection treatment developed in Section 19.5.

FIG 19.5  
FIG 19.6

KT

KT

19-28

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 879

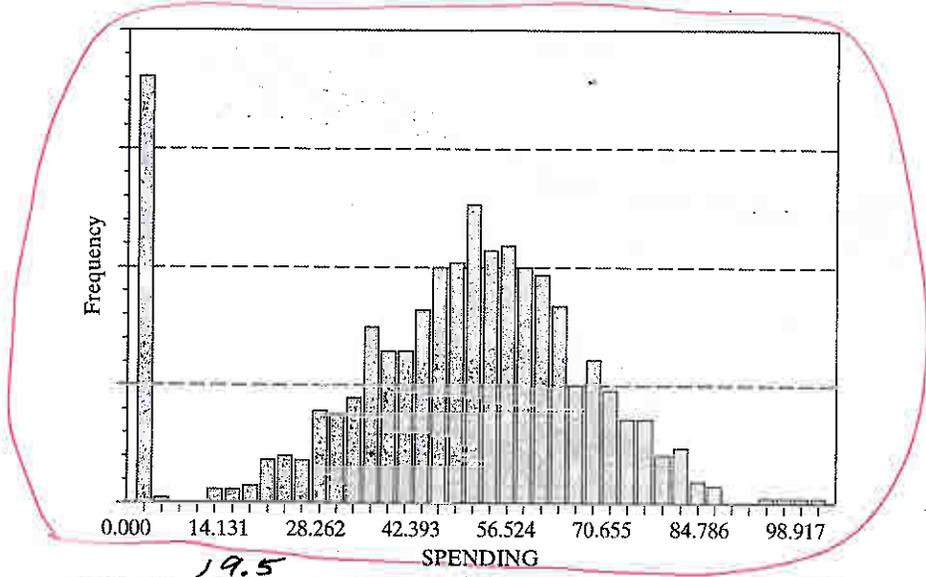
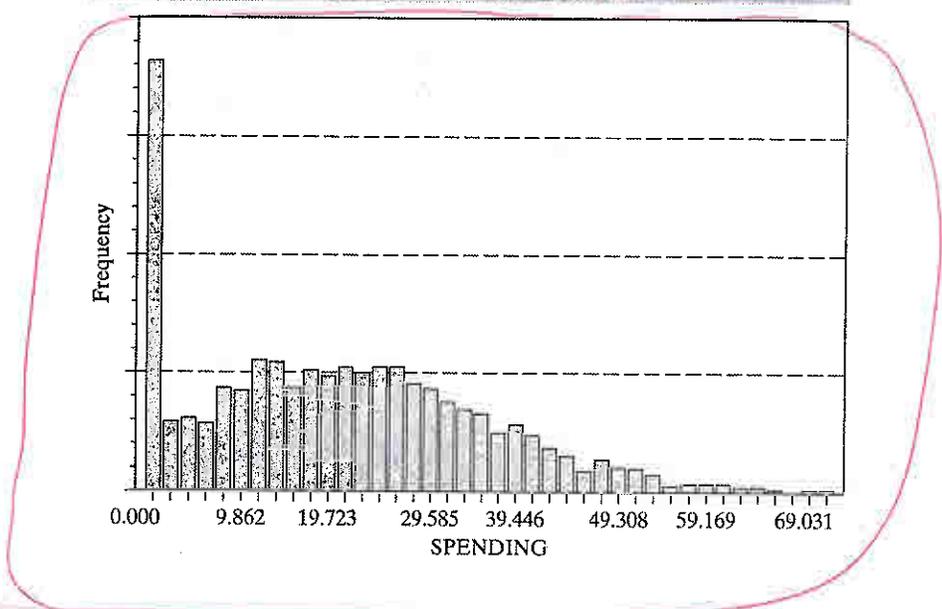


FIGURE 24.3 Hypothetical Spending Data.

Two practical issues frequently intervene at this point. First, one might well have a model in mind for the intensity (regression) equation, but none for the participation equation. This is the usual backdrop for the uses of the tobit model, which produces the considerations in the previous section. The second issue concerns the appropriateness of the truncation or censoring model to data such as those in Figure 24.3. If we consider only

19.6  
 FIGURE 24.4 Hypothetical Censored Data.



## 880 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

the nonlimit observations in Figure 24.3 the underlying distribution does not appear to be truncated at all. The truncated regression model in Section 24.2.3 fit to these data will not depart significantly from ordinary least squares [because the underlying probability in the denominator of (24-6) will equal one and the numerator will equal zero]. But, this is not the case of a tobit model forced on these same data. Forcing the model in (24-13) on data such as these will significantly distort the estimator. If all else equal, it will significantly attenuate the coefficients, the more so the larger is the proportion of limit observations in the sample. Once again, this stands as a caveat for the model builder. The tobit model is manifestly misspecified for data such as those in Figure 24.3.

**24.3.4.d Nonnormality**

Nonnormality is an especially difficult problem in this setting. It has been shown that if the underlying disturbances are not normally distributed, then the estimator based on (24-13) is inconsistent. Research is ongoing both on alternative estimators and on methods for testing for this type of misspecification.<sup>15</sup>

One approach to the estimation is to use an alternative distribution. Kalbfleisch and Prentice (2002) present a unifying treatment that includes several distributions such as the exponential, lognormal, and Weibull. (Their primary focus is on survival analysis in a medical statistics setting, which is an interesting convergence of the techniques in very different disciplines.) Of course, assuming some other specific distribution does not necessarily solve the problem and may make it worse. A preferable alternative would be to devise an estimator that is robust to changes in the distribution. Powell's (1981, 1984) least absolute deviations (LAD) estimator appears to offer some promise.<sup>16</sup> The main drawback to its use is its computational complexity. An extensive application of the LAD estimator is Melenberg and van Soest (1996). Although estimation in the nonnormal case is relatively difficult, testing for this failure of the model is worthwhile to assess the estimates obtained by the conventional methods. Among the tests that have been developed are Hausman tests, Lagrange multiplier tests [Bera and Jarque (1981, 1982), Bera, Jarque, and Lee (1982)], and **conditional moment tests** [Nelson (1981)]. The conditional moment tests are described in the next section.

To employ a Hausman test, we require an estimator that is consistent and efficient under the null hypothesis but inconsistent under the alternative—the tobit estimator with normality—and an estimator that is consistent under both hypotheses but inefficient under the null hypothesis. Thus, we will require a robust estimator of  $\beta$ , which restores the difficulties of the previous paragraph. Recent applications [e.g., Melenberg and van Soest (1996)] have used the Hausman test to compare the tobit/normal estimator with Powell's consistent, but inefficient (robust), LAD estimator. Another approach to testing is to embed the normal distribution in some other distribution and then use an LM test for the normal specification. Chesher and Irish (1987) have devised an LM test of normality in the tobit model based on **generalized residuals**. In many models, including the tobit model, the generalized residuals can be computed as the derivatives

<sup>15</sup>See Duncan (1983, 1986b), Goldberger (1983), Pagan and Vella (1989), Lee (1996), and Fernandez (1986). We will examine one of the tests more closely in the following section.

<sup>16</sup>See Duncan (1986a,b) for a symposium on the subject and Amemiya (1984). Additional references are Newey, Powell, and Walker (1990); Lee (1996); and Robinson (1988).

### Example 19.6 Two-Part Model for Extramarital Affairs

In Example 18.9, we examined Fair's (1977) Psychology Today survey data on extramarital affairs. The 601 observations in the data set are mostly zero — 451 of the 601. This feature of the data motivated the author to use a tobit model to analyze these data. In our example, we reconsidered the model, since the nonzero observations were a count, not a continuous variable. Another data set in Fair's study was the Redbook Magazine survey of 6,366 married women. Once again, the outcome variable of interest was extramarital affairs. However, in this instance, the outcome data were transformed to a measure of time spent, which, being continuous, lends itself more naturally to the tobit model we are studying here. The variables in the data set are as follows (excluding three unidentified and not used):

id	= Identification number,
C	= Constant, value = 1,
yrb	= Constructed measure of time spent in extramarital affairs,
v1	= Rating of the marriage, coded 1 to 4,
v2	= Age, in years, aggregated,
v3	= Number of years married,
v4	= Number of children, top coded at 5,
v5	= Religiosity, 1 to 4, 1 = not, 4 = very,
v6	= Education, coded 9, 12, 14, 16, 17, 20,
v7	= Wife's Occupation — Hollingshead scale
v8	= Husband's occupation — Hollingshead scale.

This is a cross section of 6,366 observations with 4,313 zeros and 2053 positive values.

Table 19.3 presents estimates of various models for yrb. The leftmost column presents the OLS estimates. The least squares estimator is inconsistent in this model. The empirical regularity that the OLS estimator appears to be biased toward zero, the more so is the smaller the proportion of limit observations. Here, the ratio, based on the tobit estimates in the second column, appears to be about 4 or 5 to 1. Likewise, the OLS estimator of  $\sigma$  appears to be greatly underestimated. This would be expected, as the OLS estimator is treating the limit observations, which have no variation in the dependent variable, as if they were nonlimit observations. The third set of results is the truncated regression estimator. In principle, the truncated regression estimator is also consistent. However, it will be less efficient as it is based on less information. In our example, this estimator seems to be quite erratic, again compared to the tobit estimator. Note, for example, the coefficient on years married, which, although it is "significant" in both cases, changes sign. The  $t$  ratio on Religiousness falls from -11.11 to -1.29 in the truncation model. The probit estimator based on  $yrb > 0$  appears next. As a rough check on the corner solution aspect of our model, we would expect the normalized tobit coefficients ( $\beta/\sigma$ ) to approximate the probit coefficients, which they appear to. However, the likelihood ratio statistic for testing the internal consistency based on the three estimated models is  $2[7804.38 - 3463.71 - 3469.58] = 1742.18$  with 9 degrees of freedom. The hypothesis of parameter constancy implied by the tobit model is rejected. The last two sets of results are for a hurdle model in which the intensity equation is fit by the two-step method.

TB  
19.3

minus

nine/

minus

Table 19.3 Estimated Censored Regression Models (t-ratios in parentheses)

	Model						
	Linear OLS	Tobit	Truncated Regression	Probit	Tobit/ $\sigma$	Hurdle Participation	Hurdle Intensity
<b>Constant</b>	3.62346 (13.63)	7.83653 (10.98)	8.89449 (2.90)	2.21010 (12.60)	1.74189	1.56419 (17.75)	4.84602 (5.87)
<b>RateMarr</b>	-0.42053 (-14.79)	-1.53071 (-20.85)	-0.44303 (-1.45)	-0.42874 (-23.40)	-0.34024	-0.42582 (-23.61)	-0.24603 (-.46)
<b>Age</b>	-0.01457 (-1.59)	-0.10514 (-4.24)	-0.22394 (-1.83)	-0.03542 (-5.87)	-0.02337		-0.01903 (-.77)
<b>YrsMarr</b>	-0.01599 (-1.62)	0.12829 (4.86)	-0.94437 (-7.27)	0.06563 (10.18)	0.02852		-0.16822 (-6.52)
<b>NumKids</b>	-0.01705 (-.57)	-0.02777 (-0.36)	-0.02280 (-0.06)	-0.00394 (-0.21)	-0.00617	0.14024 (11.55)	-0.28365 (-1.49)
<b>Religious</b>	-0.24374 (-7.83)	-0.94350 (-11.11)	-0.50490 (-1.29)	-0.22281 (-10.88)	-0.20972	-0.21466 (-10.64)	-0.05452 (-0.19)
<b>Education</b>	-0.01743 (-1.24)	-0.08598 (-2.28)	-0.06406 (-0.38)	-0.02373 (-2.60)	-0.01911		0.00338 (0.09)
<b>Wife Occ.</b>	0.06577 (2.10)	0.31284 (3.82)	0.00805 (0.02)	0.09539 (4.75)	0.06954		0.01505 (0.19)
<b>Hus. Occ.</b>	0.00405 (0.19)	0.01421 (0.26)	-0.09946 (-0.41)	0.00659 (0.49)	0.00316		-0.02911 (-0.53)
$\sigma$	2.14351	4.49887	5.46846				3.43748
<b>lnL</b>	R <sup>2</sup> =0.05479	-7804.38	-3463.71	-3469.58			

### 19.3.5 SOME ISSUES IN SPECIFICATION

Two issues that commonly arise in microeconomic data, heteroscedasticity and nonnormality, have been analyzed at length in the tobit setting.

FN 12

#### 19.3.5.a Heteroscedasticity

Maddala and Nelson (1975), Hurd (1979), Arabmazar and Schmidt (1982a,b), and Brown and Moffitt (1982) all have varying degrees of pessimism regarding how inconsistent the maximum likelihood estimator will be when heteroscedasticity occurs. Not surprisingly, the degree of censoring is the primary determinant. Unfortunately, all the analyses have been carried out in the setting of very specific models—for example, involving only a single dummy variable or one with groupwise heteroscedasticity—so the primary lesson is the very general conclusion that heteroscedasticity emerges as an obviously serious problem.

Two symposia that contain numerous results on these subjects are Blundell (1987) and Duncan (1986b). An application that explores these two issues in detail is Melenberg and van Soest (1996). Developing specification tests for the tobit model has been a popular enterprise. A sampling of the received literature includes Nelson (1981); Bera, Jarque, and Lee (1982); Chesher and Irish (1987); Chesher, Lancaster, and Irish (1985); Gouriéroux et al. (1984, 1987); Newey (1986); Rivers and Vuong (1988); Horowitz and Neumann (1989); and Pagan and Vella (1989). Newey (1985a, b) are useful references on the general subject of conditional moment testing. More general treatments of specification testing are Godfrey (1988) and Ruud (1984).

✓  
✓  
✓✓  
✓✓✓  
✓✓✓  
✓✓✓  
✓✓

19.84

TABLE 24.2 Estimates of a Tobit Model (standard errors in parentheses)

	Homoscedastic		Heteroscedastic	
	$\beta$		$\beta$	$\alpha$
Constant	-18.28 (5.10)		-4.11 (3.28)	-0.47 (0.60)
Beta	10.97 (3.61)		2.22 (2.00)	1.20 (1.81)
Nonmarket	0.65 (7.41)		0.12 (1.90)	0.08 (7.55)
Number	0.75 (5.74)		0.33 (4.50)	0.15 (4.58)
Merger	0.50 (5.90)		0.24 (3.00)	0.06 (4.17)
Option	2.56 (1.51)		2.96 (2.99)	0.83 (1.70)
In $L$	-547.30			-466.27
Sample size	200			200

One can approach the heteroscedasticity problem directly. Petersen and Waldman (1981) present the computations needed to estimate a tobit model with heteroscedasticity of several types. Replacing  $\sigma$  with  $\sigma_i$  in the log-likelihood function and including  $\sigma_i^2$  in the summations produces the needed generality. Specification of a particular model for  $\sigma_i$  provides the empirical model for estimation.

19.87

Example 24.5 Multiplicative Heteroscedasticity in the Tobit Model

Petersen and Waldman (1981) analyzed the volume of short interest in a cross section of common stocks. The regressors included a measure of the market component of heterogeneous expectations as measured by the firm's *BETA* coefficient; a company-specific measure of heterogeneous expectations, *NONMARKET*; the *NUMBER* of analysts making earnings forecasts for the company; the number of common shares to be issued for the acquisition of another firm, *MERGER*; and a dummy variable for the existence of *OPTIONS*. They report the results listed in Table 24.2 for a model in which the variance is assumed to be of the form  $\sigma_i^2 = \exp(x_i'\alpha)$ . The values in parentheses are the ratio of the coefficient to the estimated asymptotic standard error.

The effect of heteroscedasticity on the estimates is extremely large. We do note, however, a common misconception in the literature. The change in the coefficients is often misleading. The marginal effects in the heteroscedasticity model will generally be very similar to those computed from the model which assumes homoscedasticity. (The calculation is pursued in the exercises.)

A test of the hypothesis that  $\alpha = 0$  (except for the constant term) can be based on the likelihood ratio statistic. For these results, the statistic is  $-2[-547.3 - (-466.27)] = 162.06$ . This statistic has a limiting chi-squared distribution with five degrees of freedom. The sample value exceeds the critical value in the table of 11.07, so the hypothesis can be rejected.

In the preceding example, we carried out a likelihood ratio test against the hypothesis of homoscedasticity. It would be desirable to be able to carry out the test without having to estimate the unrestricted model. A Lagrange multiplier test can be used for that purpose. Consider the heteroscedastic tobit model in which we specify that

$$\sigma_i^2 = \sigma^2 [\exp(w_i'\alpha)]^2$$

19-18  
(24-15)

This model is a fairly general specification that includes many familiar ones as special cases. The null hypothesis of homoscedasticity is  $\alpha = 0$ . (We used this specification in the probit model in Section 23.4.4.b and in the linear regression model in Section 16.9.2.a) Using the BHHH estimator of the Hessian as usual, we can produce

TB  
19.4

19.3

KT

9.7.1

17.3.7

19-34

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 877

a Lagrange multiplier statistic as follows: Let  $z_i = 1$  if  $y_i$  is positive and 0 otherwise,

$$a_i = z_i \left( \frac{\varepsilon_i}{\sigma^2} \right) + (1 - z_i) \left( \frac{(-1)\lambda_i}{\sigma} \right),$$

$$b_i = z_i \left( \frac{(\varepsilon_i^2/\sigma^2 - 1)}{2\sigma^2} \right) + (1 - z_i) \left( \frac{(\mathbf{x}'_i \beta)\lambda_i}{2\sigma^3} \right), \quad \begin{matrix} 19-19 \\ (24-16) \end{matrix}$$

$$\lambda_i = \frac{\phi(\mathbf{x}'_i \beta / \sigma)}{1 - \Phi(\mathbf{x}'_i \beta / \sigma)}.$$

The data vector is  $\mathbf{g}_i = [a_i \mathbf{x}'_i, b_i, w_i]'$ . The sums are taken over all observations, and all functions involving unknown parameters ( $\varepsilon_i, \phi_i, \Phi_i, \mathbf{x}'_i \beta, \sigma, \lambda_i$ ) are evaluated at the restricted (homoscedastic) maximum likelihood estimates. Then,

$$LM = \mathbf{y}' \mathbf{G} [\mathbf{G}' \mathbf{G}]^{-1} \mathbf{G}' \mathbf{i} = nR^2 \quad \begin{matrix} 19-20 \\ (24-17) \end{matrix}$$

in the regression of a column of ones on the  $K + 1 + P$  derivatives of the log-likelihood function for the model with multiplicative heteroscedasticity, evaluated at the estimates from the restricted model. (If there were no limit observations, then it would reduce to the Breusch-Pagan statistic discussed in Section 8.5.2.) Given the maximum likelihood estimates of the tobit model coefficients, it is quite simple to compute. The statistic has a limiting chi-squared distribution with degrees of freedom equal to the number of variables in  $\mathbf{w}_i$ .

9.5.2

**24.3.4.b Misspecification of Prob[ $y_i^* > 0$ ]**

In an early study in this literature, Cragg (1971) proposed a somewhat more general model in which the probability of a limit observation is independent of the regression model for the nonlimit data. One can imagine, for instance, the decision of whether or not to purchase a car as being different from the decision of how much to spend on the car, having decided to buy one. A related problem raised by Fin and Schmidt (1984) is that in the tobit model, a variable that increases the probability of an observation being a nonlimit observation also increases the mean of the variable. They cite as an example loss due to fire in buildings. Older buildings might be more likely to have fires, so that  $\partial \text{Prob}[y_i > 0] / \partial \text{age}_i > 0$ , but, because of the greater value of newer buildings, older ones incur smaller losses when they do have fires, so that  $\partial E[y_i | y_i > 0] / \partial \text{age}_i < 0$ . This fact would require the coefficient on age to have different signs in the two functions, which is impossible in the tobit model because they are the same coefficient.

A more general model that accommodates these objections is as follows:

1. Decision equation:

$$\begin{aligned} \text{Prob}[y_i^* > 0] &= \Phi(\mathbf{x}'_i \boldsymbol{\gamma}), & z_i &= 1 \text{ if } y_i^* > 0, \\ \text{Prob}[y_i^* \leq 0] &= 1 - \Phi(\mathbf{x}'_i \boldsymbol{\gamma}), & z_i &= 0 \text{ if } y_i^* \leq 0. \end{aligned} \quad (24-18)$$

2. Regression equation for nonlimit observations:

$$E[y_i | z_i = 1] = \mathbf{x}'_i \boldsymbol{\beta} + \sigma \lambda_i,$$

according to Theorem 24.2.

This model is a combination of the truncated regression model of Section 24.2 and the univariate probit model of Section 23.3, which suggests a method of analyzing it. The tobit model of this section arises if  $\boldsymbol{\gamma} = \boldsymbol{\beta} / \sigma$ . The parameters of the regression

880 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

the nonlimit observations in Figure 24.3, the underlying distribution does not appear to be truncated at all. The truncated regression model in Section 24.2.3 fit to these data will not depart significantly from ordinary least squares [because the underlying probability in the denominator of (24-6) will equal one and the numerator will equal zero]. But, this is not the case of a tobit model forced on these same data. Forcing the model in (24-13) on data such as these will significantly distort the estimator—all else equal, it will significantly attenuate the coefficients, the more so the larger is the proportion of limit observations in the sample. Once again, this stands as a caveat for the model builder. The tobit model is manifestly misspecified for data such as those in Figure 24.3.

24.3.4.d Nonnormality

Nonnormality is an especially difficult problem in this setting. It has been shown that if the underlying disturbances are not normally distributed, then the estimator based on (24-13) is inconsistent. Research is ongoing both on alternative estimators and on methods for testing for this type of misspecification.

One approach to the estimation is to use an alternative distribution. Kalbfleisch and Prentice (2002) present a unifying treatment that includes several distributions such as the exponential, lognormal, and Weibull. (Their primary focus is on survival analysis in a medical statistics setting, which is an interesting convergence of the techniques in very different disciplines.) Of course, assuming some other specific distribution does not necessarily solve the problem and may make it worse. A preferable alternative would be to devise an estimator that is robust to changes in the distribution. Powell's (1981, 1984) least absolute deviations (LAD) estimator appears to offer some promise. The main drawback to its use is its computational complexity. An extensive application of the LAD estimator is Melenberg and van Soest (1996). Although estimation in the nonnormal case is relatively difficult, testing for this failure of the model is worthwhile to assess the estimates obtained by the conventional methods. Among the tests that have been developed are Hausman tests, Lagrange multiplier tests [Bera and Jarque (1981, 1982), Bera, Jarque, and Lee (1982)], and conditional moment tests [Nelson (1981)]. The conditional moment tests are described in the next section.

To employ a Hausman test, we require an estimator that is consistent and efficient under the null hypothesis but inconsistent under the alternative—the tobit estimator with normality—and an estimator that is consistent under both hypotheses but inefficient under the null hypothesis. Thus, we will require a robust estimator of  $\beta$ , which restores the difficulties of the previous paragraph. Recent applications [e.g., Melenberg and van Soest (1996)] have used the Hausman test to compare the tobit/normal estimator with Powell's consistent, but inefficient (robust), LAD estimator. Another approach to testing is to embed the normal distribution in some other distribution and then use an LM test for the normal specification. Chesher and Irish (1987) have devised an LM test of normality in the tobit model based on generalized residuals. In many models, including the tobit model, the generalized residuals can be computed as the derivatives

See Duncan (1983, 1986b), Goldberger (1983), Pagan and Vella (1989), Lee (1996), and Fernandez (1986). We will examine one of the tests more closely in the following section.

See Duncan (1986a,b) for a symposium on the subject and Amemiya (1984). Additional references are Newey, Powell, and Walker (1990); Lee (1996); and Robinson (1988).

Handwritten note: *the next test*

19.3.6  
#.b

19

FN 13

FN 14

16

17

15

16

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 881

of the log-densities with respect to the constant term, so

$$e_i = \frac{1}{\sigma^2} [z_i(y_i - \mathbf{x}'_i\beta) - (1 - z_i)\sigma\lambda_i],$$

where  $z_i$  is defined in (24-18) and  $\lambda_i$  is defined in (24-16). This residual is an estimate of  $\varepsilon_i$  that accounts for the censoring in the distribution. By construction,  $E[e_i | \mathbf{x}_i] = 0$ , and if the model actually does contain a constant term, then  $\sum_{i=1}^n e_i = 0$ ; this is the first of the necessary conditions for the MLE. The test is then carried out by regressing a column of ones on  $\mathbf{d}_i$ , where the elements of  $\mathbf{d}_i$  are computed as follows:

$$\begin{aligned} a_i &= (y_i - \mathbf{x}'_i\beta) / \hat{\sigma}, \\ \hat{a}_i &= -\mathbf{x}'_i\hat{\beta} / \hat{\sigma}, \lambda_i = \phi(\hat{a}_i) / \Phi(\hat{a}_i), \\ e_{i1} &= -(1 - z_i)\lambda_i + z_i a_i, \\ e_{i2} &= -(1 - z_i)\hat{a}_i \lambda_i + z_i (a_i^2 - 1), \\ e_{i3} &= -(1 - z_i) (2 + \hat{a}_i^2) \lambda_i + z_i a_i^3, \\ e_{i4} &= -(1 - z_i) (3\hat{a}_i + \hat{a}_i^3) \lambda_i + z_i (a_i^4 - 3), \\ \mathbf{d}'_i &= [e_{i1}\mathbf{x}'_i, e_{i2}, e_{i3}, e_{i4}]. \end{aligned}$$

Note that first  $K + 1$  variables in  $\mathbf{d}_i$  are the derivatives of the tobit log-likelihood. Let  $\mathbf{D}$  be the  $n \times (K + 3)$  matrix with  $i$ th row equal to  $\mathbf{d}'_i$ . Then,  $\mathbf{D} = [\mathbf{G}, \mathbf{M}]$  where the  $K + 1$  columns of  $\mathbf{G}$  are the derivatives of the tobit log-likelihood and the two columns in  $\mathbf{M}$  are the last two variables in  $\mathbf{d}_i$ . Then the chi-squared statistic is  $nR^2$ ; that is,

$$LM = \mathbf{i}'\mathbf{D}(\mathbf{D}'\mathbf{D})^{-1}\mathbf{D}'\mathbf{i}.$$

The necessary conditions that define the MLE are  $\mathbf{i}'\mathbf{G} = \mathbf{0}$ , so the first  $K + 1$  elements of  $\mathbf{i}'\mathbf{D}$  are zero. Using (A-74), then, the LM statistic becomes

$$LM = \mathbf{i}'\mathbf{M}[\mathbf{M}'\mathbf{M} - \mathbf{M}'\mathbf{G}(\mathbf{G}'\mathbf{G})^{-1}\mathbf{G}'\mathbf{M}]^{-1}\mathbf{M}'\mathbf{i},$$

which is a chi-squared statistic with two degrees of freedom. Note the similarity to (24-17), where a test for homoscedasticity is carried out by the same method. As emerges so often in this framework, the test of the distribution actually focuses on the skewness and kurtosis of the residuals. Pagan and Vella (1989) and Ruud (1984) have developed similar tests for several specification errors.<sup>17</sup>

19.3.7  
6  
29/12/15

24.4 PANEL DATA APPLICATIONS

17 Extension of the familiar panel data results to the tobit model parallel the probit model, with the attendant problems. The random effects or random parameters models discussed in Chapter 23 can be adapted to the censored regression model using simulation or quadrature. The same reservations with respect to the orthogonality of the effects and the regressors will apply here, as will the applicability of the Mundlak (1978) correction to accommodate it.

move to previous page

<sup>17</sup>Developing specification tests for the tobit model has been a popular enterprise. A sampling of the received literature includes Nelson (1981); Bera, Jarque, and Lee (1982); Chesher and Irish (1987); Chesher, Lancaster, and Irish (1985); Courrieroux et al. (1984, 1987); Newey (1986); Rivers and Vuong (1988); Horowitz and Neumann (1989); and Pagan and Vella (1989). Newey (1985a, b) are useful references on the general subject of conditional moment testing. More general treatments of specification testing are Godfrey (1988) and Ruud (1984).

## 882 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

Most of the attention in the theoretical literature on panel data methods for the tobit model has been focused on fixed effects. The departure point would be the maximum likelihood estimator for the static fixed effects model.

$$y_{it}^* = \alpha_i + x'_{it}\beta + \varepsilon_{it}, \varepsilon_{it} \sim N[0, \sigma^2],$$

$$y_{it} = \text{Max}(0, y_{it}^*).$$

However, there are no firm theoretical results on the behavior of the MLE in this model. Intuition might suggest, based on the findings for the binary probit model, that the MLE would be biased in the same fashion, away from zero. Perhaps surprisingly, the results in Greene (2004) persistently found that not to be the case in a variety of model specifications. Rather, the incidental parameters, such as it is, manifests in a downward bias in the estimator of  $\sigma$ , not an upward (or downward) bias in the MLE of  $\beta$ . However, this is less surprising when the tobit estimator is juxtaposed with the MLE in the linear regression model with fixed effects. In that model, the MLE is the within-groups (LSDV) estimator which is unbiased and consistent. But, the ML estimator of the disturbance variance in the linear regression model is  $\hat{\sigma}_{LSDV}^2 = \text{eLSDV}/(nT)$ , which is biased downward by a factor of  $(T-1)/T$ . [This is the result found in the original source on the incidental parameters problem, Neyman and Scott (1948).] So, what evidence there is suggests that unconditional estimation of the tobit model behaves essentially like that for the linear regression model. That does not settle the problem, however; if the evidence is correct, then it implies that although consistent estimation of  $\beta$  is possible, appropriate statistical inference is not. The bias in the estimation of  $\sigma$  shows up in any estimator of the asymptotic covariance of the MLE of  $\beta$ .

Unfortunately, there is no conditional estimator of  $\beta$  for the tobit (or truncated regression) model. First differencing or taking group mean deviations does not preserve the model. Because the latent variable is censored before observation, these transformations are not meaningful. Some progress has been made on theoretical, semiparametric estimators for this model. See, for example, Honoré and Kyriazidou (2000) for a survey. Much of the theoretical development has also been directed at dynamic models where the benign result of the previous paragraph (such as it is) is lost once again. Arellano (2001) contains some general results. Hahn and Kuersteiner (2004) have characterized the bias of the MLE, and suggested methods of reducing the bias of the estimators in dynamic binary choice and censored regression models.

## 24.5 SAMPLE SELECTION

The topic of sample selection, or **incidental truncation**, has been the subject of an enormous recent literature, both theoretical and applied.<sup>18</sup> This analysis combines both of the previous topics.

<sup>18</sup>A large proportion of the analysis in this framework has been in the area of labor economics. See, for example, Vella (1998), which is an extensive survey for practitioners. The results, however, have been applied in many other fields, including, for example, long series of stock market returns by financial economists ("survivorship bias") and medical treatment and response in long-term studies by clinical researchers ("attrition bias"). The four surveys noted in the introduction to this chapter provide fairly extensive, although far from exhaustive, lists of the studies. Some studies that comment on methodological issues are Heckman (1990), Manski (1989, 1990, 1992), and Newey, Powell, and Walker (1990).

five

## 19.4 MODELS FOR DURATION

The leading application of the censoring models we examined in Section 19.3 is models for durations and events. We consider the time until some kind of transition as the duration, and the transition, itself, as the event. The length of a spell of unemployment (until rehire or exit from the market), the duration of a strike, the amount of time until a patient ends a health related spell in connection with a disease or operation, and the length of time between origination and termination (via prepayment, default or some other mechanism) of a mortgage are all examples of durations and transitions. The role that censoring plays in these scenarios is that in almost all cases in which we as analysts study duration data, some or even many of the spells we observe do not end in transitions. For example, in studying the lengths of unemployment spells, many of the individuals in the sample may still be unemployed at the time the study ends  $\frac{1}{M}$  the analyst observes (or believes) that the spell will end some time after the observation window closes. These data on spell lengths are, by construction, censored. Models of duration will generally account explicitly for censoring of the duration data.

This section is concerned with models of duration. In some aspects, the regression-like models we have studied, such as the discrete choice models, are the appropriate tools. As in the previous two chapters, however, the models are nonlinear, and the familiar regression methods are not appropriate. Most of this analysis focuses on maximum likelihood estimators. In modeling duration, although an underlying regression model is, in fact, at work, it is generally not the conditional mean function that is of interest. More likely, as we will explore below, the objects of estimation are certain probabilities of events, for example in the conditional probability of a transition in a given interval given that the spell has lasted up to the point of interest. These are known as "hazard models"  $\frac{1}{M}$  the probability is labelled the hazard function  $\frac{1}{M}$  and are a central focus of this type of analysis. *next/*

## CHAPTER 25 ♦ Models for Event Counts and Duration 931

19  
TABLE 25.7 Estimates of a Zero-Inflated Poisson Model

Variable	Poisson Regression		Logit Splitting Model		Marginal Effects		
	1.1 Estimate	Standard Error	Estimate	Standard Error	ZIP	Tobit (0)	Tobit (0, 4)
Constant	1.27	0.439	-1.85	0.664	—	—	—
Age	-0.00422	0.0122	0.0397	0.0190	-0.0216	-0.0420	-0.0218
Years	0.0331	0.0231	-0.0981	0.0318	0.0702	0.130	0.0654
Religion	-0.0909	0.0721	0.306	0.0951	-0.210	-0.394	-0.199
Occupation	0.0205	0.0441	0.0677	0.0607	0.0467	0.0762	0.0399
Happiness	-0.817	0.0666	0.458	0.0949	-0.273	-0.534	-0.271

Because, as expected, the ancillary model for the zeros accounted for the overdispersion in the data, the negative binomial model degenerated to the Poisson form.) Finally, the marginal effects,  $\delta = \partial E[y | x] / \partial x$ , are shown in Table 25.7 for three models: the ZIP model, Fair's original tobit model, and the tobit model estimated with the doubly censored count. The estimates for the ZIP model are considerably lower than those for Fair's tobit model. When the tobit model is reestimated with the censoring on the right, however, the resulting marginal effects are reasonably close to those from the ZIP model, although uniformly smaller. (This result may be from not building the censoring into the ZIP model, a refinement that would be relatively straightforward.)

19 We conclude that the original tobit model provided only a fair approximation to the marginal effects produced by (we contend) the more appropriate specification of the Poisson model. But the approximation became much better when the data were recoded and treated as censored. Figure 25.1 also shows the predictions from the ZIP model (narrow bars). As might be expected, it provides a much better prediction of the dependent variable. (The integer values of the conditional mean function for the tobit model were roughly evenly split between zeros and ones, whereas the doubly censored model always predicted  $y = 0$ .) Surprisingly, the treatment of the highest observations does greatly affect the outcome. If the ZIP model is fit to the original uncensored data, then the vector of marginal effects is  $\delta = [-0.0586, 0.2446, -0.692, 0.115, -0.787]$ , which is extremely large. Thus, perhaps more analysis is called for—the ZIP model can be further improved, and one might reconsider the hurdle model—but we have tortured Fair's data enough. Further exploration is left for the reader.

## 19.4.1 19 25.6 MODELS FOR DURATION DATA 17

Intuition might suggest that the longer a strike persists, the more likely it is that it will end within, say, the next week. Or is it? It seems equally plausible to suggest that the longer a strike has lasted, the more difficult must be the problems that led to it in the first place, and hence the *less* likely it is that it will end in the next short time interval.

17 There are a large number of highly technical articles on this topic but relatively few accessible sources for the uninitiated. A particularly useful introductory survey is Kiefer (1988), upon which we have drawn heavily for this section. Other useful sources are Kalbfleisch and Prentice (2002), Heckman and Singer (1984a), Lancaster (1990), and Florens, Fougere, and Mouchart (1996), and Cameron and Trivedi;

(2005, chapters 17-19).

## 932 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

A similar kind of reasoning could be applied to spells of unemployment or the interval between conceptions. In each of these cases, it is not only the duration of the event, per se, that is interesting, but also the likelihood that the event will end in "the next period" given that it has lasted as long as it has.

Analysis of the length of *time until failure* has interested engineers for decades. For example, the models discussed in this section were applied to the durability of electric and electronic components long before economists discovered their usefulness. Likewise, the analysis of *survival times*—for example, the length of survival after the onset of a disease or after an operation such as a heart transplant—has long been a staple of biomedical research. Social scientists have recently applied the same body of techniques to strike duration, length of unemployment spells, intervals between conception, time until business failure, length of time between arrests, length of time from purchase until a warranty claim is made, intervals between purchases, and so on.

This section will give a brief introduction to the econometric analysis of duration data. As usual, we will restrict our attention to a few straightforward, relatively uncomplicated techniques and applications, primarily to introduce terms and concepts. The reader can then wade into the literature to find the extensions and variations. We will concentrate primarily on what are known as **parametric models**. These apply familiar inference techniques and provide a convenient departure point. Alternative approaches are considered at the end of the discussion.

19.4.2

## 25.3 + DURATION DATA

The variable of interest in the analysis of duration is the length of time that elapses from the beginning of some event either until its end or until the measurement is taken, which may precede termination. Observations will typically consist of a cross section of durations,  $t_1, t_2, \dots, t_n$ . The process being observed may have begun at different points in calendar time for the different individuals in the sample. For example, the strike duration data examined in Example 25.4 are drawn from nine different years.

19.8

Censoring is a pervasive and usually unavoidable problem in the analysis of duration data. The common cause is that the measurement is made while the process is ongoing. An obvious example can be drawn from medical research. Consider analyzing the survival times of heart transplant patients. Although the beginning times may be known with precision, at the time of the measurement, observations on any individuals who are still alive are necessarily censored. Likewise, samples of spells of unemployment drawn from surveys will probably include some individuals who are still unemployed at the time the survey is taken. For these individuals, duration, or survival, is at least the observed  $t_i$ , but not equal to it. Estimation must account for the censored nature of the data for the same reasons as considered in Section 18.3. The consequences of ignoring censoring in duration data are similar to those that arise in regression analysis.

18 19.3

In a conventional regression model that characterizes the conditional mean and variance of a distribution, the regressors can be taken as fixed characteristics at the point in time or for the individual for which the measurement is taken. When measuring duration, the observation is implicitly on a process that has been under way for an interval of time from zero to  $t$ . If the analysis is conditioned on a set of covariates (the counterparts to regressors)  $x_t$ , then the duration is implicitly a function of the entire

19-41

CHAPTER 25 ♦ Models for Event Counts and Duration 933

time path of the variable  $x(t)$ ,  $t = (0, t)$ , which may have changed during the interval. For example, the observed duration of employment in a job may be a function of the individual's rank in the firm. But their rank may have changed several times between the time they were hired and when the observation was made. As such, observed rank at the end of the job tenure is not necessarily a complete description of the individual's rank while they were employed. Likewise, marital status, family size, and amount of education are all variables that can change during the duration of unemployment and that one would like to account for in the duration model. The treatment of time-varying covariates is a considerable complication. \* 18 16

FN 120

KT

KT

19.4.3

25.4.2

A REGRESSION-LIKE APPROACH: PARAMETRIC MODELS OF DURATION

We will use the term spell as a catchall for the different duration variables we might measure. Spell length is represented by the random variable  $T$ . A simple approach to duration analysis would be to apply regression analysis to the sample of observed spells. By this device, we could characterize the expected duration, perhaps conditioned on a set of covariates whose values were measured at the end of the period. We could also assume that conditioned on an  $x$  that has remained fixed from  $T=0$  to  $T=t$ ,  $t$  has a normal distribution, as we commonly do in regression. We could then characterize the probability distribution of observed duration times. But, normality turns out not to be particularly attractive in this setting for a number of reasons, not least of which is that duration is positive by construction, while a normally distributed variable can take negative values. (Lognormality turns out to be a palatable alternative, but it is only one among a long list of candidates.)

19.4.3.a

~~25.4.2~~ Theoretical Background

Suppose that the random variable  $T$  has a continuous probability distribution  $f(t)$ , where  $t$  is a realization of  $T$ . The cumulative probability is

$$F(t) = \int_0^t f(s) ds = \text{Prob}(T \leq t).$$

We will usually be more interested in the probability that the spell is of length at least  $t$ , which is given by the survival function, KT

$$S(t) = 1 - F(t) = \text{Prob}(T \geq t).$$

Consider the question raised in the introduction: Given that the spell has lasted until time  $t$ , what is the probability that it will end in the next short interval of time, say,  $\Delta t$ ? It is

$$l(t, \Delta t) = \text{Prob}(t \leq T \leq t + \Delta t | T \geq t).$$

A useful function for characterizing this aspect of the distribution is the hazard rate, KT

$$\lambda(t) = \lim_{\Delta t \rightarrow 0} \frac{\text{Prob}(t \leq T \leq t + \Delta t | T \geq t)}{\Delta t} = \lim_{\Delta t \rightarrow 0} \frac{F(t + \Delta t) - F(t)}{\Delta t S(t)} = \frac{f(t)}{S(t)}.$$

18 See Petersen (1986) for one approach to this problem.

## 934 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

Roughly, the hazard rate is the rate at which spells are completed after duration  $t$ , given that they last at least until  $t$ . As such, the hazard function gives an answer to our original question.

The hazard function, the density, the CDF and the survival function are all related. The hazard function is

$$\lambda(t) = \frac{-d \ln S(t)}{dt},$$

so

$$f(t) = S(t)\lambda(t).$$

Another useful function is the integrated hazard function

$$\Lambda(t) = \int_0^t \lambda(s) ds,$$

for which

$$S(t) = e^{-\Lambda(t)},$$

so

$$\Lambda(t) = -\ln S(t).$$

The integrated hazard function is generalized residual in this setting. [See Chesher and Irish (1987) and Example 26.4] 19.8

19.4.3. b

~~19.4.3. b~~ **Models of the Hazard Function**

For present purposes, the hazard function is more interesting than the survival rate or the density. Based on the previous results, one might consider modeling the hazard function itself, rather than, say, modeling the survival function then obtaining the density and the hazard. For example, the base case for many analyses is a hazard rate that does not vary over time. That is,  $\lambda(t)$  is a constant  $\lambda$ . This is characteristic of a process that has no memory; the conditional probability of "failure" in a given short interval is the same regardless of when the observation is made. Thus, and

$$\lambda(t) = \lambda.$$

From the earlier definition, we obtain the simple differential equation,

$$\frac{-d \ln S(t)}{dt} = \lambda.$$

The solution is

$$\ln S(t) = k - \lambda t,$$

or

$$S(t) = Ke^{-\lambda t},$$

where  $K$  is the constant of integration. The terminal condition that  $S(0) = 1$  implies that  $K = 1$ , and the solution is

$$S(t) = e^{-\lambda t}.$$

19-43

CHAPTER 25 ♦ Models for Event Counts and Duration 935

This solution is the <sup>(KT)</sup> exponential distribution, which has been used to model the time until failure of electronic components. Estimation of  $\lambda$  is simple, because with an exponential distribution,  $E[t] = 1/\lambda$ . The maximum likelihood estimator of  $\lambda$  would be the reciprocal of the sample mean.

A natural extension might be to model the hazard rate as a linear function,  $\lambda(t) = \alpha + \beta t$ . Then  $\Lambda(t) = \alpha t + \frac{1}{2}\beta t^2$  and  $f(t) = \lambda(t)S(t) = \lambda(t) \exp[-\Lambda(t)]$ . To avoid a negative hazard function, one might depart from  $\lambda(t) = \exp[g(t, \theta)]$ , where  $\theta$  is a vector of parameters to be estimated. With an observed sample of durations, estimation of  $\alpha$  and  $\beta$  is, at least in principle, a straightforward problem in maximum likelihood. [Kennan (1985) used a similar approach.]

A distribution whose hazard function slopes upward is said to have <sup>(KT)</sup> positive duration dependence. For such distributions, the likelihood of failure at time  $t$ , conditional upon duration up to time  $t$ , is increasing in  $t$ . The opposite case is that of decreasing hazard or <sup>(KT)</sup> negative duration dependence. Our question in the introduction about whether the strike is more or less likely to end at time  $t$  given that it has lasted until time  $t$  can be framed in terms of positive or negative duration dependence. The assumed distribution has a considerable bearing on the answer. If one is unsure at the outset of the analysis whether the data can be characterized by positive or negative duration dependence, then it is counterproductive to assume a distribution that displays one characteristic or the other over the entire range of  $t$ . Thus, the exponential distribution and our suggested extension could be problematic. The literature contains a cornucopia of choices for duration models: normal, inverse normal [inverse Gaussian; see Lancaster (1990)], lognormal,  $F$ , gamma, Weibull (which is a popular choice), and many others.<sup>19</sup> To illustrate the differences, we will examine a few of the simpler ones. Table 25.8 lists the hazard functions and survival functions for four commonly used distributions. Each involves two parameters, a location parameter,  $\lambda$ , and a scale parameter,  $p$ . [Note that in the benchmark case of the exponential distribution,  $\lambda$  is the hazard function. In all other cases, the hazard function is a function of  $\lambda$ ,  $p$ , and, where there is duration dependence,  $t$  as well. Different authors, e.g., Kiefer (1988), use different parameterizations of these models. We follow the convention of Kalbfleisch and Prentice (2002).]

All these are distributions for a nonnegative random variable. Their hazard functions display very different behaviors, as can be seen in Figure 25.2. The hazard function for the exponential distribution is constant, that for the Weibull is monotonically increasing or decreasing depending on  $p$ , and the hazards for lognormal and loglogistic

AV: Term "exponential" is not in Chap. list

(KT)  
(KT)

FN 17  
TB 195

for example

FIG 19.7

19.5

AV: OK to spell out "eg."?

19.7

19.5  
TABLE 25.8 Survival Distributions

Distribution	Hazard Function, $\lambda(t)$	Survival Function, $S(t)$
Exponential	$\lambda$	$S(t) = e^{-\lambda t}$
Weibull	$\lambda p(\lambda t)^{p-1}$	$S(t) = e^{-(\lambda t)^p}$
Lognormal	$f(t) = (p/t)\phi[p \ln(\lambda t)]$ [ $\ln t$ is normally distributed with mean $-\ln \lambda$ and standard deviation $1/p$ .]	$S(t) = \Phi[-p \ln(\lambda t)]$
Loglogistic	$\lambda(t) = \lambda p(\lambda t)^{p-1} / [1 + (\lambda t)^p]$ [ $\ln t$ has a logistic distribution with mean $-\ln \lambda$ and variance $\pi^2 / (3p^2)$ .]	$S(t) = 1 / [1 + (\lambda t)^p]$

19 Three sources that contain numerous specifications are Kalbfleisch and Prentice (2002), Cox and Oakes (1985), and Lancaster (1990).

19-44

936 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

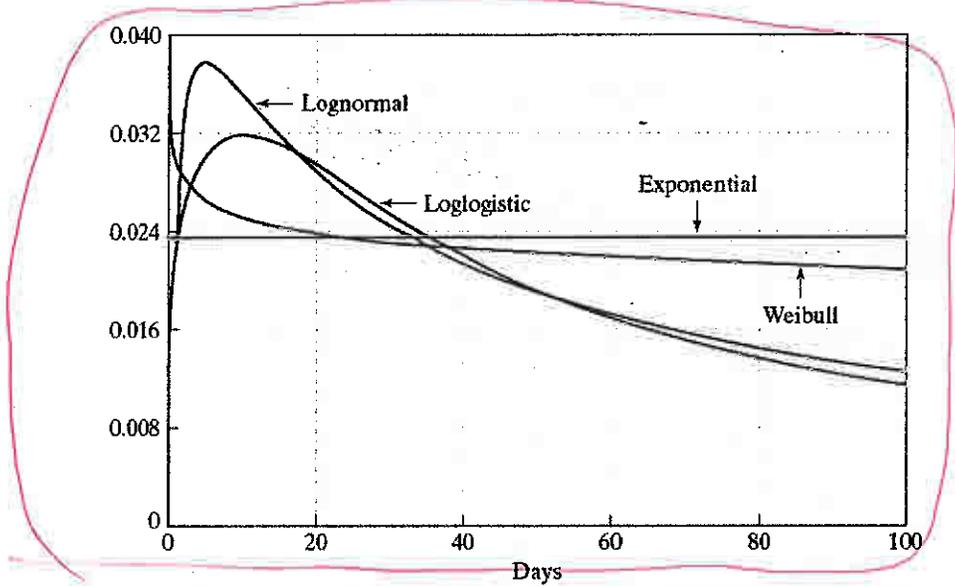


FIGURE 25.2 Parametric Hazard Functions.

19.7

distributions first increase and then decrease. Which among these or the many alternatives is likely to be best in any application is uncertain.

19.4.3.c ~~19.4.3.c~~ 25.6.2.c Maximum Likelihood Estimation

The parameters  $\lambda$  and  $p$  of these models can be estimated by maximum likelihood. For observed duration data,  $t_1, t_2, \dots, t_n$ , the log-likelihood function can be formulated and maximized in the ways we have become familiar with in earlier chapters. Censored observations can be incorporated as in Section 24.3 for the tobit model. [See (24-13).] As such,

19.3 19-13

$$\ln L(\theta) = \sum_{\text{uncensored observations}} \ln f(t|\theta) + \sum_{\text{censored observations}} \ln S(t|\theta),$$

where  $\theta = (\lambda, p)$ . For some distributions, it is convenient to formulate the log-likelihood function in terms of  $f(t) = \lambda(t)S(t)$  so that

$$\ln L = \sum_{\text{uncensored observations}} \ln \lambda(t|\theta) + \sum_{\text{all observations}} \ln S(t|\theta).$$

Inference about the parameters can be done in the usual way. Either the BHHH estimator or actual second derivatives can be used to estimate asymptotic standard errors for the estimates. The transformation  $w = p(\ln t + \ln \lambda)$  for these distributions greatly facilitates maximum likelihood estimation. For example, for the Weibull model, by defining  $w = p(\ln t + \ln \lambda)$ , we obtain the very simple density  $f(w) = \exp[w - \exp(w)]$  and

19-45

CHAPTER 25 ♦ Models for Event Counts and Duration 937

FN 18

survival function  $S(w) = \exp(-\exp(w))$ .<sup>18 20</sup> Therefore, by using  $\ln t$  instead of  $t$ , we greatly simplify the log-likelihood function. Details for these and several other distributions may be found in Kalbfleisch and Prentice (2002, pp. 68-70). The Weibull distribution is examined in detail in the next section.

19.4.3. d ~~19.4.3. d~~ Exogenous Variables

One limitation of the models given earlier is that external factors are not given a role in the survival distribution. The addition of "covariates" to duration models is fairly straightforward, although the interpretation of the coefficients in the model is less so. Consider, for example, the Weibull model. (The extension to other distributions will be similar.) Let

$$\lambda_i = e^{-\mathbf{x}'_i \boldsymbol{\beta}}$$

where  $\mathbf{x}_i$  is a constant term and a set of variables that are assumed not to change from time  $T=0$  until the "failure time,"  $T=t_i$ . Making  $\lambda_i$  a function of a set of regressors is equivalent to changing the units of measurement on the time axis. For this reason, these models are sometimes called **accelerated failure time models**. Note as well that in all the models listed (and generally), the regressors do not bear on the question of duration dependence, which is a function of  $p$ .

Let  $\sigma = 1/p$  and let  $\delta_i = 1$  if the spell is completed and  $\delta_i = 0$  if it is censored. As before, let

$$w_i = p \ln(\lambda_i t_i) = \frac{(\ln t_i - \mathbf{x}'_i \boldsymbol{\beta})}{\sigma}$$

and denote the density and survival functions  $f(w_i)$  and  $S(w_i)$ . The observed random variable is

$$\ln t_i = \sigma w_i + \mathbf{x}'_i \boldsymbol{\beta}$$

The Jacobian of the transformation from  $w_i$  to  $\ln t_i$  is  $dw_i/d \ln t_i = 1/\sigma$ , so the density and survival functions for  $\ln t_i$  are

$$f(\ln t_i | \mathbf{x}_i, \boldsymbol{\beta}, \sigma) = \frac{1}{\sigma} f\left(\frac{\ln t_i - \mathbf{x}'_i \boldsymbol{\beta}}{\sigma}\right), \text{ and } S(\ln t_i | \mathbf{x}_i, \boldsymbol{\beta}, \sigma) = S\left(\frac{\ln t_i - \mathbf{x}'_i \boldsymbol{\beta}}{\sigma}\right).$$

The log-likelihood for the observed data is

$$\ln L(\boldsymbol{\beta}, \sigma | \text{data}) = \sum_{i=1}^n [\delta_i \ln f(\ln t_i | \mathbf{x}_i, \boldsymbol{\beta}, \sigma) + (1 - \delta_i) \ln S(\ln t_i | \mathbf{x}_i, \boldsymbol{\beta}, \sigma)].$$

For the **Weibull model**, for example (see footnote N)<sup>KT</sup>

$$f(w_i) = \exp(w_i - e^{w_i}),$$

and

$$S(w_i) = \exp(-e^{w_i}).$$

FN 18

18 20

The transformation is  $\exp(w) = (\lambda t)^p$  so  $t = (1/\lambda)[\exp(w)]^{1/p}$ . The Jacobian of the transformation is  $dt/dw = [\exp(w)]^{1/p} / (\lambda p)$ . The density in Table 25.8 is  $\lambda p [\exp(w)]^{-(1/p)-1} [\exp(-\exp(w))]$ . Multiplying by the Jacobian produces the result,  $f(w) = \exp[w - \exp(w)]$ . The survival function is the antiderivative,  $[\exp(-\exp(w))]$ .

☺

Av: Term "Weibull model" is not in chap. list

☺

19.5  
Av: I8 x-ref to Table 19.5 OK?

938 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

Making the transformation to  $\ln t_i$  and collecting terms reduces the log-likelihood to

$$\ln L(\beta, \sigma | \text{data}) = \sum_i \left[ \delta_i \left( \frac{\ln t_i - x_i' \beta}{\sigma} - \ln \sigma \right) - \exp \left( \frac{\ln t_i - x_i' \beta}{\sigma} \right) \right]$$

19.5

(Many other distributions, including the others in Table 25.8, simplify in the same way. The exponential model is obtained by setting  $\sigma$  to one.) The derivatives can be equated to zero using the methods described in Section E.3. The individual terms can also be used to form the BHHH estimator of the asymptotic covariance matrix for the estimator. The Hessian is also simple to derive, so Newton's method could be used instead.

FN 19

FN 20

Note that the hazard function generally depends on  $t$ ,  $p$ , and  $x$ . The sign of an estimated coefficient suggests the direction of the effect of the variable on the hazard function when the hazard is monotonic. But in those cases, such as the loglogistic, in which the hazard is nonmonotonic, even this may be ambiguous. The magnitudes of the effects may also be difficult to interpret in terms of the hazard function. In a few cases, we do get a regression-like interpretation. In the Weibull and exponential models,  $E[t | x_i] = \exp(x_i' \beta) \Gamma[(1/p) + 1]$ , whereas for the lognormal and loglogistic models,  $E[\ln t | x_i] = x_i' \beta$ . In these cases,  $\beta_k$  is the derivative (or a multiple of the derivative) of this conditional mean. For some other distributions, the conditional median of  $t$  is easily obtained. Numerous cases are discussed by Kiefer (1988), Kalbfleisch and Prentice (2002), and Lancaster (1990).

19.4.3.e ~~25.6.2~~ Heterogeneity

The problem of heterogeneity in duration models can be viewed essentially as the result of an incomplete specification. Individual specific covariates are intended to incorporate observation specific effects. But if the model specification is incomplete and if systematic individual differences in the distribution remain after the observed effects are accounted for, then inference based on the improperly specified model is likely to be problematic. We have already encountered several settings in which the possibility of heterogeneity mandated a change in the model specification; the fixed and random effects regression, logit, and probit models all incorporate observation-specific effects. Indeed, all the failures of the linear regression model discussed in the preceding chapters can be interpreted as a consequence of heterogeneity arising from an incomplete specification.

There are a number of ways of extending duration models to account for heterogeneity. The strictly nonparametric approach of the Kaplan-Meier estimator (see Section 25.6.3) is largely immune to the problem, but it is also rather limited in how much information can be culled from it. One direct approach is to model heterogeneity in the parametric model. Suppose that we posit a survival function conditioned on the individual specific effect  $v_i$ . We treat the survival function as  $S(t_i | v_i)$ . Then add to that a model for the unobserved heterogeneity  $f(v_i)$ . (Note that this is a counterpart to the incorporation of a disturbance in a regression model and follows the same procedures

19.4.4

19.3.2

21 Note that the log-likelihood function has the same form as that for the tobit model in Section 24.3. By just reinterpreting the nonlimit observations in a tobit setting, we can, therefore, use this framework to apply a wide range of distributions to the tobit model. [See Greene (1995a) and references given therein.]  
 22 See Kalbfleisch and Prentice (2002) for numerous other examples.

CHAPTER 25 ♦ Models for Event Counts and Duration 939

that we used in the Poisson model with random effects.) Then

$$S(t) = E_v[S(t|v)] = \int_v S(t|v) f(v) dv. \quad 21, 23$$

FN 21

The gamma distribution is frequently used for this purpose.<sup>23</sup> Consider, for example, using this device to incorporate heterogeneity into the Weibull model we used earlier. As is typical, we assume that  $v$  has a gamma distribution with mean 1 and variance  $\theta = 1/k$ . Then

$$f(v) = \frac{k^k}{\Gamma(k)} e^{-kv} v^{k-1},$$

and

$$S(t|v) = e^{-(v\lambda t)^p}.$$

After a bit of manipulation, we obtain the unconditional distribution,

$$S(t) = \int_0^\infty S(t|v) f(v) dv = [1 + \theta(\lambda t)^p]^{-1/\theta}.$$

FN 22

The limiting value, with  $\theta = 0$ , is the **Weibull survival model**, so  $\theta = 0$  corresponds to  $\text{Var}[v] = 0$ , or no heterogeneity.<sup>25</sup> The hazard function for this model is

$$\lambda(t) = \lambda p (\lambda t)^{p-1} [S(t)]^\theta, \quad 24, 25$$

which shows the relationship to the Weibull model.

This approach is common in parametric modeling of heterogeneity. In an important paper on this subject, Heckman and Singer (1984b) argued that this approach tends to overparameterize the survival distribution and can lead to rather serious errors in inference. They gave some dramatic examples to make the point. They also expressed some concern that researchers tend to choose the distribution of heterogeneity more on the basis of mathematical convenience than on any sensible economic basis.

19.4.4 ~~19.4.3~~ NONPARAMETRIC AND SEMIPARAMETRIC APPROACHES

The parametric models are attractive for their simplicity. But by imposing as much structure on the data as they do, the models may distort the estimated hazard rates. It may be that a more accurate representation can be obtained by imposing fewer restrictions.

The Kaplan-Meier (1958) **product limit estimator** is a strictly empirical, nonparametric approach to survival and hazard function estimation. Assume that the observations on duration are sorted in ascending order so that  $t_1 \leq t_2$  and so on and, for now, that no observations are censored. Suppose as well that there are  $K$  distinct survival times in the data, denoted  $T_k$ ;  $K$  will equal  $n$  unless there are ties. Let  $n_k$  denote

23

<sup>23</sup> See, for example, Hausman, Hall, and Griliches (1984), who use it to incorporate heterogeneity in the Poisson regression model. The application is developed in Section 15.3.

24

<sup>24</sup> For the strike data analyzed in Figure 25.2, the maximum likelihood estimate of  $\theta$  is 0.0004, which suggests that at least in the context of the Weibull model, latent heterogeneity does not appear to be a feature of these data.

18.4.4  
19

19.7

19-48

940 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

the number of individuals whose observed duration is at least  $T_k$ . The set of individuals whose duration is at least  $T_k$  is called the **risk set** at this duration. (We borrow, once again, from biostatistics, where the risk set is those individuals still "at risk" at time  $T_k$ ). Thus,  $n_k$  is the size of the risk set at time  $T_k$ . Let  $h_k$  denote the number of observed spells completed at time  $T_k$ . A strictly empirical estimate of the survivor function would be

$$\hat{S}(T_k) = \prod_{i=1}^k \frac{n_i - h_i}{n_i} = \frac{n_i - h_i}{n_1}$$

The estimator of the hazard rate is

$$\hat{\lambda}(T_k) = \frac{h_k}{n_k}$$

19-21  
(25-14)

Corrections are necessary for observations that are censored. Lawless (1982), Kalbfleisch and Prentice (2002), Kiefer (1988), and Greene (1995a) give details. Susin (2001) points out a fundamental ambiguity in this calculation (one which he argues appears in the 1958 source). The estimator in (25-14) is not a "rate" as such, as the width of the time window is undefined, and could be very different at different points in the chain of calculations. Because many intervals, particularly those late in the observation period, might have zeros, the failure to acknowledge these intervals should impart an upward bias to the estimator. His proposed alternative computes the counterpart to (25-14) over a mesh of defined intervals as follows:

19-21

19-21

$$\hat{\lambda}(I_a^b) = \frac{\sum_{j=a}^b h_j}{\sum_{j=a}^b n_j b_j}$$

where the interval is from  $t = a$  to  $t = b$ ,  $h_j$  is the number of failures in each period in this interval,  $n_j$  is the number of individuals at risk in that period and  $b_j$  is the width of the period. Thus, an interval  $(a, b)$  is likely to include several "periods."

(KT) Cox's (1972) approach to the **proportional hazard** model is another popular, **semi-parametric** method of analyzing the effect of covariates on the hazard rate. The model specifies that

$$\lambda(t_i) = \exp(x_i' \beta) \lambda_0(t_i)$$

The function  $\lambda_0$  is the "baseline" hazard, which is the individual heterogeneity. In principle, this hazard is a parameter for each observation that must be estimated. Cox's **partial likelihood** estimator provides a method of estimating  $\beta$  without requiring estimation of  $\lambda_0$ . The estimator is somewhat similar to Chamberlain's estimator for the logit model with panel data in that a conditioning operation is used to remove the heterogeneity. (See Section 25.2.) Suppose that the sample contains  $K$  distinct exit times,  $T_1, \dots, T_K$ .

17.4.4

For any time  $T_k$ , the risk set, denoted  $R_k$ , is all individuals whose exit time is at least  $T_k$ . The risk set is defined with respect to any moment in time  $T$  as the set of individuals who have not yet exited just prior to that time. For every individual  $i$  in risk set  $R_k$ ,  $t_i \geq T_k$ . The probability that an individual exits at time  $T_k$  given that exactly one

AV: KT is chap list is "semiparametric model" OK?

CHAPTER 25 ♦ Models for Event Counts and Duration 941

*individual exits at this time* (which is the counterpart to the conditioning in the binary logit model in Chapter 23) is

17 
$$\text{Prob}[t_i = T_k | \text{risk set}_k] = \frac{e^{x_i \beta}}{\sum_{j \in R_k} e^{x_j \beta}}$$

Thus, the conditioning sweeps out the baseline hazard functions. For the simplest case in which exactly one individual exits at each distinct exit time and there are no censored observations, the partial log-likelihood is

$$\ln L = \sum_{k=1}^K \left[ x_k' \beta - \ln \sum_{j \in R_k} e^{x_j \beta} \right]$$

If  $m_k$  individuals exit at time  $T_k$ , then the contribution to the log-likelihood is the sum of the terms for each of these individuals.

The proportional hazard model is a common choice for modeling durations because it is a reasonable compromise between the Kaplan-Meier estimator and the possibly excessively structured parametric models. Hausman and Han (1990) and Meyer (1988), among others, have devised other, "semiparametric" specifications for hazard models.

19.8

**Example 25.4 Survival Models for Strike Duration**

19.32

The strike duration data given in Kennan (1985, pp. 14-16) have become a familiar standard for the demonstration of hazard models. Appendix Table F25.8 lists the durations, in days, of 62 strikes that commenced in June of the years 1968 to 1976. Each involved at least 1,000 workers and began at the expiration or reopening of a contract. Kennan reported the actual duration. In his survey, Kiefer (1985), using the same observations, censored the data at 80 days to demonstrate the effects of censoring. We have kept the data in their original form; the interested reader is referred to Kiefer for further analysis of the censoring problem.

23  
25

Parameter estimates for the four duration models are given in Table 25.9. The estimate of the median of the survival distribution is obtained by solving the equation  $S(t) = 0.5$ . For example, for the Weibull model,

$$S(M) = 0.5 = \exp[-(\lambda M)^p],$$

19.6

or

$$M = [(\ln 2)^{1/p}] / \lambda.$$

For the exponential model,  $p = 1$ . For the lognormal and loglogistic models,  $M = 1/\lambda$ . The delta method is then used to estimate the standard error of this function of the parameter estimates. (See Section 4.8.4.) All these distributions are skewed to the right. As such,  $E[t]$  is

FN 23  
TB 19.6

19.6  
**TABLE 25.9** Estimated Duration Models (estimated standard errors in parentheses)

4.4.4

	$\lambda$	$p$	Median Duration
Exponential	0.02344 (0.00298)	1.00000 (0.00000)	29.571 (3.522)
Weibull	0.02439 (0.00354)	0.92083 (0.11086)	27.543 (3.997)
Loglogistic	0.04153 (0.00707)	1.33148 (0.17201)	24.079 (4.102)
Lognormal	0.04514 (0.00806)	0.77206 (0.08865)	22.152 (3.954)

23  
25 Our statistical results are nearly the same as Kiefer's despite the censoring.

## 942 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

greater than the median. For the exponential and Weibull models,  $E[t] = [1/\lambda]\Gamma[(1/p) + 1]$ ; for the normal,  $E[t] = (1/\lambda)[\exp(1/p^2)]^{1/2}$ . The implied hazard functions are shown in Figure 25.2.

The variable  $x$  reported with the strike duration data is a measure of unanticipated aggregate industrial production net of seasonal and trend components. It is computed as the residual in a regression of the log of industrial production in manufacturing on time, time squared, and monthly dummy variables. With the industrial production variable included as a covariate, the estimated Weibull model is

$$-\ln \lambda = 3.7772 - 9.3515x, \quad p = 1.00288,$$

$$(0.1394) (2.973) \quad (0.1217),$$

median strike length = 27.35(3.667) days,  $E[t] = 39.83$  days.

Note that the Weibull model is now almost identical to the exponential model ( $p = 1$ ). Because the hazard conditioned on  $x$  is approximately equal to  $\lambda$ , it follows that the hazard function is increasing in "unexpected" industrial production. A 1 percent increase in  $x$  leads to a 9.35 percent increase in  $\lambda$ , which because  $p \approx 1$  translates into a 9.35 percent decrease in the median strike length or about 2.6 days. (Note that  $M = \ln 2/\lambda$ .)

The proportional hazard model does not have a constant term. (The baseline hazard is an individual specific constant.) The estimate of  $\beta$  is  $-9.0726$ , with an estimated standard error of 3.225. This is very similar to the estimate obtained for the Weibull model.

## 28.7 SUMMARY AND CONCLUSIONS

This chapter has surveyed models for events. We analyze these data in two forms. Models for "count data" are appropriate for describing numbers of events, such as hospital visits, arrivals of messages at a network node, and so on. The Poisson model is the central pillar of this area, although its limitations usually motivate researchers to consider alternative specifications. The count data models are essentially nonlinear regressions, but it is more fruitful to do the modeling in terms of the probabilities of events rather than as a model of a conditional mean. We considered basic cases, as well as extensions that accommodate features of data such as censoring, truncation, zero inflation, and unobserved heterogeneity. The second type of event history data that we considered are durations. It is useful to think of duration, or survival data, as the measurement of time between transitions or changes of state. We examined three modeling approaches that correspond to the description in Chapter 14, nonparametric (survival tables), semiparametric (the proportional hazard models), and parametric (various forms such as the Weibull model).

### Key Terms and Concepts

- Accelerated failure time model
- Censored variable
- Censoring
- Conditional moment test
- Count data
- Delta method
- Deviance
- Duration dependence
- Duration model
- Exponential model
- Exposure
- Generalized residual
- Hazard function
- Hazard rate
- Heterogeneity
- Hurdle model
- Integrated hazard function
- Lagrange multiplier test
- Loglinear model
- Marginal effects
- Negative binomial model
- Negative duration dependence
- Negbin 1 form
- Negbin 2 form
- Negbin  $P$  model
- Nonnested models
- Overdispersion
- Parametric model
- Partial likelihood
- Poisson regression model

882 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

Most of the attention in the theoretical literature on panel data methods for the tobit model has been focused on fixed effects. The departure point would be the maximum likelihood estimator for the static fixed effects model,

$$y_{it}^* = \alpha_i + x_{it}'\beta + \varepsilon_{it}, \varepsilon_{it} \sim N[0, \sigma^2],$$

$$y_{it} = \text{Max}(0, y_{it}^*).$$

However, there are no firm theoretical results on the behavior of the MLE in this model. Intuition might suggest, based on the findings for the binary probit model, that the MLE would be biased in the same fashion, away from zero. Perhaps surprisingly, the results in Greene (2004) persistently found that not to be the case in a variety of model specifications. Rather, the incidental parameters, such as it is, manifests in a downward bias in the estimator of  $\sigma$ , not an upward (or downward) bias in the MLE of  $\beta$ . However, this is less surprising when the tobit estimator is juxtaposed with the MLE in the linear regression model with fixed effects. In that model, the MLE is the within-groups (LSDV) estimator which is unbiased and consistent. But, the ML estimator of the disturbance variance in the linear regression model is  $e'_{\text{LSDV}}e_{\text{LSDV}}/(nT)$ , which is biased downward by a factor of  $(T-1)/T$ . [This is the result found in the original source on the incidental parameters problem, Neyman and Scott (1948).] So, what evidence there is suggests that unconditional estimation of the tobit model behaves essentially like that for the linear regression model. That does not settle the problem, however: if the evidence is correct, then it implies that although consistent estimation of  $\beta$  is possible, appropriate statistical inference is not. The bias in the estimation of  $\sigma$  shows up in any estimator of the asymptotic covariance of the MLE of  $\beta$ .

Unfortunately, there is no conditional estimator of  $\beta$  for the tobit (or truncated regression) model. First differencing or taking group mean deviations does not preserve the model. Because the latent variable is censored before observation, these transformations are not meaningful. Some progress has been made on theoretical, **semiparametric estimators** for this model. See, for example, Honoré and Kyriazidou (2000) for a survey. Much of the theoretical development has also been directed at dynamic models where the benign result of the previous paragraph (such as it is) is lost once again. Arellano (2001) contains some general results. Hahn and Kuersteiner (2004) have characterized the bias of the MLE, and suggested methods of reducing the bias of the estimators in dynamic binary choice and censored regression models.

19.5 ~~24.5~~ INCIDENTAL TRUNCATION AND SAMPLE SELECTION

The topic of sample selection, or **incidental truncation**, has been the subject of an enormous recent literature, both theoretical and applied.<sup>18</sup> This analysis combines both of the previous topics.

A large proportion of the analysis in this framework has been in the area of labor economics. See, for example, Vella (1998), which is an extensive survey for practitioners. The results, however, have been applied in many other fields, including, for example, long series of stock market returns by financial economists ("survivorship bias") and medical treatment and response in long-term studies by clinical researchers ("attrition bias"). ~~The four surveys noted in the introduction to this chapter provide fairly extensive, although far from exhaustive, coverage of the literature.~~ Some studies that comment on methodological issues are Heckman (1990), Manski (1989, 1990, 1992), and Newey, Powell, and Walker (1990).

FN 24

24 26

KT

24 26

Handwritten signature

19-52

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 883

19.9

Example 24.6. Incidental Truncation 19.2

In the high-income survey discussed in Example 24.2, respondents were also included in the survey if their net worth, not including their homes, was at least \$500,000. Suppose that the survey of incomes was based only on people whose net worth was at least \$500,000. This selection is a form of truncation, but not quite the same as in Section 24.2. This selection criterion does not necessarily exclude individuals whose incomes at the time might be quite low. Still, one would expect that, on average, individuals with a high net worth would have a high income as well. Thus, the average income in this subpopulation would in all likelihood also be misleading as an indication of the income of the typical American. The data in such a survey would be nonrandomly selected or incidentally truncated. 19.2

Econometric studies of nonrandom sampling have analyzed the deleterious effects of sample selection on the properties of conventional estimators such as least squares; have produced a variety of alternative estimation techniques; and, in the process, have yielded a rich crop of empirical models. In some cases, the analysis has led to a reinterpretation of earlier results.

19.5.1

~~24.5.1~~ 19.5.1 INCIDENTAL TRUNCATION IN A BIVARIATE DISTRIBUTION

Suppose that  $y$  and  $z$  have a bivariate distribution with correlation  $\rho$ . We are interested in the distribution of  $y$  given that  $z$  exceeds a particular value. Intuition suggests that if  $y$  and  $z$  are positively correlated, then the truncation of  $z$  should push the distribution of  $y$  to the right. As before, we are interested in (1) the form of the incidentally truncated distribution and (2) the mean and variance of the incidentally truncated random variable. Because it has dominated the empirical literature, we will focus first on the bivariate normal distribution. 24

The truncated joint density of  $y$  and  $z$  is

$$f(y, z | z > a) = \frac{f(y, z)}{\text{Prob}(z > a)}$$

To obtain the incidentally truncated marginal density for  $y$ , we would then integrate  $z$  out of this expression. The moments of the incidentally truncated normal distribution are given in Theorem 24.5. 27 25

FN 25

<sup>19.5</sup>  
~~24.5~~ **THEOREM 24.5** Moments of the Incidentally Truncated Bivariate Normal Distribution

If  $y$  and  $z$  have a bivariate normal distribution with means  $\mu_y$  and  $\mu_z$ , standard deviations  $\sigma_y$  and  $\sigma_z$ , and correlation  $\rho$ , then

$$E[y | z > a] = \mu_y + \rho\sigma_y\lambda(\alpha_z),$$

$$\text{Var}[y | z > a] = \sigma_y^2[1 - \rho^2\delta(\alpha_z)],$$

19.5.2  
(24-19)

where

$$\alpha_z = (a - \mu_z)/\sigma_z, \lambda(\alpha_z) = \phi(\alpha_z)/[1 - \Phi(\alpha_z)], \text{ and } \delta(\alpha_z) = \lambda(\alpha_z)[\lambda(\alpha_z) - \alpha_z].$$

25 27

<sup>19</sup>We will reconsider the issue of the normality assumption in Section 24.5.5.

<sup>20</sup>Much more general forms of the result that apply to multivariate distributions are given in Johnson and Kotz (1974). See also Maddala (1983, pp. 266–267).

884 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

Note that the expressions involving  $z$  are analogous to the moments of the truncated distribution of  $x$  given in Theorem 24.2. If the truncation is  $z < a$ , then we make the replacement  $\lambda(\alpha_z) = -\phi(\alpha_z)/\Phi(\alpha_z)$ .

As expected, the truncated mean is pushed in the direction of the correlation if the truncation is from below and in the opposite direction if it is from above. In addition, the incidental truncation reduces the variance, because both  $\delta(\alpha)$  and  $\rho^2$  are between zero and one.

19.5.2 ~~24.5.2~~ REGRESSION IN A MODEL OF SELECTION

To motivate a regression model that corresponds to the results in Theorem 24.5, we consider the following example.

19.10 Example 24.7 A Model of Labor Supply

A simple model of female labor supply that has been examined in many studies consists of two equations:

FN  
26

1. *Wage equation.* The difference between a person's *market wage*, what she could command in the labor market, and her *reservation wage*, the wage rate necessary to make her choose to participate in the labor market, is a function of characteristics such as age and education as well as, for example, number of children and where a person lives.
2. *Hours equation.* The desired number of labor hours supplied depends on the wage, home characteristics such as whether there are small children present, marital status, and so on.

The problem of truncation surfaces when we consider that the second equation describes desired hours, but an actual figure is observed only if the individual is working. (In most such studies, only a *participation equation*, that is, whether hours are positive or zero, is observable.) We infer from this that the market wage exceeds the reservation wage. Thus, the hours variable in the second equation is incidentally truncated.

To put the preceding examples in a general framework, let the equation that determines the sample selection be

$$z_i^* = w_i' \gamma + u_i,$$

and let the equation of primary interest be

$$y_i = x_i' \beta + \varepsilon_i.$$

The sample rule is that  $y_i$  is observed only when  $z_i^*$  is greater than zero. Suppose as well that  $\varepsilon_i$  and  $u_i$  have a bivariate normal distribution with zero means and correlation  $\rho$ . Then we may insert these in Theorem 24.5 to obtain the model that applies to the observations in our sample:

19.5

$$\begin{aligned} E[y_i | y_i \text{ is observed}] &= E[y_i | z_i^* > 0] \quad \leftarrow -\frac{w_i' \gamma}{\sigma_u} \\ &= E[y_i | u_i > -w_i' \gamma] \\ &= x_i' \beta + E[\varepsilon_i | u_i > -w_i' \gamma] \\ &= x_i' \beta + \rho \sigma_\varepsilon \lambda_i(\alpha_u) \\ &= x_i' \beta + \beta_\lambda \lambda_i(\alpha_u), \end{aligned}$$

26  
28

\*Sec. for example, Heckman (1976). This strand of literature begins with an exchange by Gronau (1974) and Lewis (1974).

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 885

where  $\alpha_u = -w_i' \gamma / \sigma_u$  and  $\lambda(\alpha_u) = \phi(w_i' \gamma / \sigma_u) / \Phi(w_i' \gamma / \sigma_u)$ . So,

$$y_i | z_i^* > 0 = E[y_i | z_i^* > 0] + v_i = x_i' \beta + \beta_\lambda \lambda_i(\alpha_u) + v_i$$

Least squares regression using the observed data—for instance, OLS regression of hours on its determinants, using only data for women who are working—produces inconsistent estimates of  $\beta$ . Once again, we can view the problem as an omitted variable. Least squares regression of  $y$  on  $x$  and  $\lambda$  would be a consistent estimator, but if  $\lambda$  is omitted, then the specification error of an omitted variable is committed. Finally, note that the second part of Theorem 24.5 implies that even if  $\lambda_i$  were observed, then least squares would be inefficient. The disturbance  $v_i$  is heteroscedastic.

19.5 KT

The marginal effect of the regressors on  $y_i$  in the observed sample consists of two components. There is the direct effect on the mean of  $y_i$ , which is  $\beta$ . In addition, for a particular independent variable, if it appears in the probability that  $z_i^*$  is positive, then it will influence  $y_i$  through its presence in  $\lambda_i$ . The full effect of changes in a regressor that appears in both  $x_i$  and  $w_i$  on  $y$  is

$$\frac{\partial E[y_i | z_i^* > 0]}{\partial x_{ik}} = \beta_k - \gamma_k \left( \frac{\rho \sigma_\varepsilon}{\sigma_u} \right) \delta_i(\alpha_u),$$

FN 27

where

$$\delta_i = \lambda_i^2 - \alpha_i \lambda_i$$

Suppose that  $\rho$  is positive and  $E[y_i]$  is greater when  $z_i^*$  is positive than when it is negative. Because  $0 < \delta_i < 1$ , the additional term serves to reduce the marginal effect. The change in the probability affects the mean of  $y_i$  in that the mean in the group  $z_i^* > 0$  is higher. The second term in the derivative compensates for this effect, leaving only the marginal effect of a change given that  $z_i^* > 0$  to begin with. Consider Example 24.9, and suppose that education affects both the probability of migration and the income in either state. If we suppose that the income of migrants is higher than that of otherwise identical people who do not migrate, then the marginal effect of education has two parts, one due to its influence in increasing the probability of the individual's entering a higher-income group and one due to its influence on income within the group. As such, the coefficient on education in the regression overstates the marginal effect of the education of migrants and understates it for nonmigrants. The sizes of the various parts depend on the setting. It is quite possible that the magnitude, sign, and statistical significance of the effect might all be different from those of the estimate of  $\beta$ , a point that appears frequently to be overlooked in empirical studies.

19.12

In most cases, the selection variable  $z^*$  is not observed. Rather, we observe only its sign. To consider our two examples, we typically observe only whether a woman is working or not working or whether an individual migrated or not. We can infer the sign of  $z^*$ , but not its magnitude, from such information. Because there is no information on the scale of  $z^*$ , the disturbance variance in the selection equation cannot be estimated. (We encountered this problem in Chapter 23 in connection with the probit model.)

27

<sup>22</sup>We have reversed the sign of  $\alpha_u$  in (24-19) because  $a = 0$ , and  $\alpha = w_i' \gamma / \sigma_u$  is somewhat more convenient. Also, as such,  $\partial \lambda / \partial \alpha = -\delta$ .

Theorem 19.5

886 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

Thus, we reformulate the model as follows:

selection mechanism:  $z_i^* = \mathbf{w}_i' \boldsymbol{\gamma} + u_i, z_i = 1$  if  $z_i^* > 0$  and 0 otherwise;

$\text{Prob}(z_i = 1 | \mathbf{w}_i) = \Phi(\mathbf{w}_i' \boldsymbol{\gamma})$ , and

$\text{Prob}(z_i = 0 | \mathbf{w}_i) = 1 - \Phi(\mathbf{w}_i' \boldsymbol{\gamma})$ .

*DATA (24-20) (19-22)*

regression model:  $y_i = \mathbf{x}_i' \boldsymbol{\beta} + \varepsilon_i$  observed only if  $z_i = 1$ ,

$(u_i, \varepsilon_i) \sim \text{bivariate normal } [0, 0, 1, \sigma_\varepsilon, \rho]$ .

Suppose that, as in many of these studies,  $z_i$  and  $\mathbf{w}_i$  are observed for a random sample of individuals but  $y_i$  is observed only when  $z_i = 1$ . This model is precisely the one we examined earlier, with

$E[y_i | z_i = 1, \mathbf{x}_i, \mathbf{w}_i] = \mathbf{x}_i' \boldsymbol{\beta} + \rho \sigma_\varepsilon \lambda(\mathbf{w}_i' \boldsymbol{\gamma})$ .

19.5.3 TWO-STEP AND MAXIMUM LIKELIHOOD ESTIMATION

The parameters of the sample selection model can be estimated by maximum likelihood.<sup>28</sup> However, Heckman's (1979) two-step estimation procedure is usually used instead. Heckman's method is as follows:<sup>29</sup>

*28 FN 28*  
*29 FN 29*

1. Estimate the probit equation by maximum likelihood to obtain estimates of  $\boldsymbol{\gamma}$ . For each observation in the selected sample, compute  $\hat{\lambda}_i = \phi(\mathbf{w}_i' \hat{\boldsymbol{\gamma}}) / \Phi(\mathbf{w}_i' \hat{\boldsymbol{\gamma}})$  and  $\hat{\delta}_i = \hat{\lambda}_i (\hat{\lambda}_i + \mathbf{w}_i' \hat{\boldsymbol{\gamma}})$ .
2. Estimate  $\boldsymbol{\beta}$  and  $\beta_\lambda = \rho \sigma_\varepsilon$  by least squares regression of  $y$  on  $\mathbf{x}$  and  $\hat{\lambda}$ .

It is possible also to construct consistent estimators of the individual parameters  $\rho$  and  $\sigma_\varepsilon$ . At each observation, the true conditional variance of the disturbance would be

$\sigma_i^2 = \sigma_\varepsilon^2 (1 - \rho^2 \delta_i)$ .

The average conditional variance for the sample would converge to

$\text{plim} \frac{1}{n} \sum_{i=1}^n \sigma_i^2 = \sigma_\varepsilon^2 (1 - \rho^2 \bar{\delta})$ ,

which is what is estimated by the least squares residual variance  $\mathbf{e}'\mathbf{e}/n$ . For the square of the coefficient on  $\lambda$ , we have

$\text{plim} b_\lambda^2 = \rho^2 \sigma_\varepsilon^2$ ,

whereas based on the probit results we have

$\text{plim} \frac{1}{n} \sum_{i=1}^n \hat{\delta}_i = \bar{\delta}$ .

We can then obtain a consistent estimator of  $\sigma_\varepsilon^2$  using

$\hat{\sigma}_\varepsilon^2 = \frac{1}{n} \mathbf{e}'\mathbf{e} + \hat{\delta} b_\lambda^2$ .

<sup>28</sup>See Greene (1995a).

<sup>29</sup>Perhaps in a mimicry of the "tobit" estimator described earlier, this procedure has come to be known as the "Heckit" estimator.

19-56

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 887

Finally, an estimator of  $\rho^2$  is

$$\hat{\rho}^2 = \frac{b_\lambda^2}{\hat{\sigma}_\epsilon^2}$$

(19-23)

FN 30

which provides a complete set of estimators of the model's parameters.

To test hypotheses, an estimate of the asymptotic covariance matrix of  $[b', b_\lambda]$  is needed. We have two problems to contend with. First, we can see in Theorem 24.5 that the disturbance term in

$$(y_i | z_i = 1, x_i, w_i) = x_i' \beta + \rho \sigma_\epsilon \lambda_i + v_i$$

19-24  
(24-21)

is heteroscedastic;

$$\text{Var}[v_i | z_i = 1, x_i, w_i] = \sigma_\epsilon^2 (1 - \rho^2 \delta_i).$$

Second, there are unknown parameters in  $\lambda_i$ . Suppose that we assume for the moment that  $\lambda_i$  and  $\delta_i$  are known (i.e., we do not have to estimate  $\gamma$ ). For convenience, let  $x_i^* = [x_i, \lambda_i]$ , and let  $b^*$  be the least squares coefficient vector in the regression of  $y$  on  $x^*$  in the selected data. Then, using the appropriate form of the variance of ordinary least squares in a heteroscedastic model from Chapter 8, we would have to estimate

$$\begin{aligned} \text{Var}[b^*] &= \sigma_\epsilon^2 [X_*' X_*]^{-1} \left[ \sum_{i=1}^n (1 - \rho^2 \delta_i) x_i^* x_i^{*'} \right] [X_*' X_*]^{-1} \\ &= \sigma_\epsilon^2 [X_*' X_*]^{-1} [X_*' (I - \rho^2 \Delta) X_*] [X_*' X_*]^{-1}, \end{aligned}$$

where  $I - \rho^2 \Delta$  is a diagonal matrix with  $(1 - \rho^2 \delta_i)$  on the diagonal. Without any other complications, this result could be computed fairly easily using  $X$ , the sample estimates of  $\sigma_\epsilon^2$  and  $\rho^2$ , and the assumed known values of  $\lambda_i$  and  $\delta_i$ .

The parameters in  $\gamma$  do have to be estimated using the probit equation. Rewrite (24-21) as

19-24

$$(y_i | z_i = 1, x_i, w_i) = x_i' \beta + \beta_\lambda \hat{\lambda}_i + v_i - \beta_\lambda (\hat{\lambda}_i - \lambda_i).$$

In this form, we see that in the preceding expression we have ignored both an additional source of variation in the compound disturbance and correlation across observations: the same estimate of  $\gamma$  is used to compute  $\hat{\lambda}_i$  for every observation. Heckman has shown that the earlier covariance matrix can be appropriately corrected by adding a term inside the brackets,

$$Q = \hat{\rho}^2 (X_*' \hat{\Delta} W) \text{Est. Asy. Var}[\hat{\gamma}] (W' \hat{\Delta} X_*) = \hat{\rho}^2 \hat{F} \hat{V} \hat{F}',$$

where  $\hat{V} = \text{Est. Asy. Var}[\hat{\gamma}]$ , the estimator of the asymptotic covariance of the probit coefficients. Any of the estimators in (23-22) to (23-24) may be used to compute  $\hat{V}$ . The complete expression is

FN 31

$$\text{Est. Asy. Var}[b, b_\lambda] = \hat{\sigma}_\epsilon^2 [X_*' X_*]^{-1} [X_*' (I - \hat{\rho}^2 \hat{\Delta}) X_* + Q] [X_*' X_*]^{-1}.$$

<sup>31</sup>Note that  $\hat{\rho}^2$  is not a sample correlation and, as such, is not limited to  $[0, 1]$ . See Greene (1981) for discussion.

<sup>32</sup>This matrix formulation is derived in Greene (1981). Note that the Murphy and Topel (1985) results for two-step estimators given in Theorem 11.3 would apply here as well. Asymptotically, this method would give the same answer. The Heckman formulation has become standard in the literature.

The sample selection model can also be estimated by maximum likelihood. The full log likelihood function for the data is built up from

and

Prob(selection)  $\times$  density | selection for observations with  $z_i = 1$ ,

and

Prob(non-selection) for observations with  $z_i = 0$ .

Combining the parts produces the full log likelihood function,

$$\ln L = \sum_{z=1} \ln \left[ \frac{\exp\left(-\frac{1}{2}\varepsilon_i^2 / \sigma_\varepsilon^2\right)}{\sigma_\varepsilon \sqrt{2\pi}} \Phi\left(\frac{\rho\varepsilon_i / \sigma_\varepsilon + \mathbf{w}'\boldsymbol{\gamma}}{\sqrt{1-\rho^2}}\right) \right] + \sum_{z=0} [1 - \ln \Phi(\mathbf{w}'\boldsymbol{\gamma})],$$

where  $\varepsilon_i = y_i - \mathbf{x}_i'\boldsymbol{\beta}$ . Note, the FIML estimator with its assumption of bivariate normality is not less robust than the two step estimator. because the latter also requires bivariate normality to form the conditional mean for the regression.

Two virtues of the FIML estimator will be the greater efficiency brought by using the likelihood function rather than the method of moments and, second, the estimation of  $\rho$  subject to the constraint  $-1 < \rho < 1$ . (This is typically done by reparameterizing the model in terms of the monotonic inverse hyperbolic tangent,  $\tau = (1/2)\ln[(1+\rho)/(1-\rho)] = \text{atanh}(\rho)$ . The transformed parameter,  $\tau$ , is unrestricted. The inverse transformation is  $\rho = [\exp(2\tau)-1]/[\exp(2\tau)+1]$  which is bounded between zero and one.) One possible drawback (it might be argued) could be the complexity of the likelihood function that would make estimation more difficult than the two step estimator. However, the MLE for the selection model appears as a built in procedure in modern software such as *Stata* and *NLOGIT*, and it is straightforward to implement in *Gauss* and *MatLab*, so this might be a moot point. Surprisingly, the MLE is by far less common than the two step estimator in the received applications. The estimation of  $\rho$  is the difficult part of the estimation process (this is often the case). It is quite common for the method of moments estimator and the FIML estimator to be very different. Our application in Example 19.11 is a case. Perhaps surprisingly so, the moment based estimator of  $\rho$  in (19-23) is not bounded by zero and one. [See Greene (1981).] This would seem to recommend the MLE.

The fully parametric bivariate normality assumption of the model has been viewed as a potential drawback. However, relatively little progress has been made on devising informative semi- and nonparametric estimators - see, for one example, Gallant and Nychka (1987). The obstacle here is that, ultimately, the model hangs on a parameterization of the correlation of the unobservables in the two equations. So, method of moment estimators or kernel based estimators must still incorporate this feature of a bivariate distribution. Some results have been obtained using the method of copula functions. [See Smith (2003, 2005) and Trivedi and Zimmer (2007).]

888 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

19.7

TABLE 24.3 Estimated Selection Corrected Wage Equation

	Two-Step		Maximum Likelihood		Least Squares	
	Estimate	Std. Err.	Estimate	Std. Err.	Estimate	Std. Err.
$\beta_1$	-0.971	(2.06)	-1.963	(1.684)	-2.56	(0.929)
$\beta_2$	0.021	(0.0625)	0.0279	(0.0756)	0.0325	(0.0616)
$\beta_3$	0.000137	(0.00188)	-0.0001	(0.00234)	-0.000260	(0.00184)
$\beta_4$	0.417	(0.100)	0.457	(0.0964)	0.481	(0.0669)
$\beta_5$	0.444	(0.316)	0.447	(0.427)	0.318	(0.449)
$(\rho\sigma)$	-1.098	(1.266)				
$\rho$	-0.343		-0.132	(0.224)	0.000	
$\sigma$	3.200		3.108	(0.0837)	3.111	

19.11

Example 24.3 Female Labor Supply

17.1 17.8

Examples 23.1 and 23.5 proposed a labor force participation model for a sample of 753 married women in a sample analyzed by Mroz (1987). The data set contains wage and hours information for the 428 women who participated in the formal market ( $LFP = 1$ ). Following Mroz, we suppose that for these 428 individuals, the offered wage exceeded the reservation wage and, moreover, the unobserved effects in the two wage equations are correlated. As such, a wage equation based on the market data should account for the sample selection problem. We specify a simple wage model:

$$Wage = \beta_1 + \beta_2 \text{Exper} + \beta_3 \text{Exper}^2 + \beta_4 \text{Education} + \beta_5 \text{City} + \varepsilon$$

18  
19.7

where Exper is labor market experience and City is a dummy variable indicating that the individual lived in a large urban area. Maximum likelihood, Heckman two-step, and ordinary least squares estimates of the wage equation are shown in Table 24.3. The maximum likelihood estimates are FIML estimates; the labor force participation equation is reestimated at the same time. Only the parameters of the wage equation are shown below. Note as well that the two-step estimator estimates the single coefficient on  $\lambda_i$  and the structural parameters  $\sigma$  and  $\rho$  are deduced by the method of moments. The maximum likelihood estimator computes estimates of these parameters directly. [Details on maximum likelihood estimation may be found in Maddala (1983).]

19.7

next

The differences between the two-step and maximum likelihood estimates in Table 24.3 are surprisingly large. The difference is even more striking in the marginal effects. The effect for education is estimated as 0.417 + 0.0641 for the two-step estimators and 0.149 in total for the maximum likelihood estimates. For the kids variable, the marginal effect is -0.293 for the two-step estimates and only -0.0113 for the MLEs. Surprisingly, the direct test for a selection effect in the maximum likelihood estimates, a nonzero  $\rho$ , fails to reject the hypothesis that  $\rho$  equals zero.

19.7

0.480

-1.1003

In some settings, the selection process is a nonrandom sorting of individuals into two or more groups. The mover-stayer model in the next example is a familiar case.

19.12

Example 24.4 A Mover-Stayer Model for Migration

The model of migration analyzed by Nakosteen and Zimmer (1980) fits into the framework described in this section. The equations of the model are

net benefit of moving:  $M_i^* = w_i^* \gamma + u_i$

income if moves:  $I_{i1} = X_{i1}' \beta_1 + \varepsilon_{i1}$

income if stays:  $I_{i0} = X_{i0}' \beta_0 + \varepsilon_{i0}$

One component of the net benefit is the market wage individuals could achieve if they move, compared with what they could obtain if they stay. Therefore, among the determinants of

19.8  
TABLE 24.4 Estimated Earnings Equations

	Migration	Migrant Earnings	Nonmigrant Earnings
Constant	-1.509	9.041	8.593
SE	-0.708 (-5.72)	-4.104 (-9.54)	-4.161 (-57.71)
$\Delta EMP$	-1.488 (-2.60)	—	—
$\Delta PCI$	1.455 (3.14)	—	—
Age	-0.008 (-5.29)	—	—
Race	-0.065 (-1.17)	—	—
Sex	-0.082 (-2.14)	—	—
$\Delta SIC$	0.948 (24.15)	-0.790 (-2.24)	-0.927 (-9.35)
$\lambda$	—	0.212 (0.50)	0.863 (2.84)

the net benefit are factors that also affect the income received in either place. An analysis of income in a sample of migrants must account for the incidental truncation of the mover's income on a positive net benefit. Likewise, the income of the stayer is incidentally truncated on a nonpositive net benefit. The model implies an income after moving for all observations, but we observe it only for those who actually do move. Nakosteen and Zimmer (1980) applied the selectivity model to a sample of 9,223 individuals with data for 2 years (1971 and 1973) sampled from the Social Security Administration's Continuous Work History Sample. Over the period, 1,078 individuals migrated and the remaining 8,145 did not. The independent variables in the migration equation were as follows:

- SE = self-employment dummy variable; 1 if yes,
- $\Delta EMP$  = rate of growth of state employment,
- $\Delta PCI$  = growth of state per capita income,
- x = age, race (nonwhite = 1), sex (female = 1),
- $\Delta SIC$  = 1 if individual changes industry.

The earnings equations included  $\Delta SIC$  and SE. The authors reported the results given in Table 24.4. The figures in parentheses are asymptotic *t* ratios.

19.8

24.5.4 REGRESSION ANALYSIS OF TREATMENT EFFECTS

The basic model of selectivity outlined earlier has been extended in an impressive variety of directions.<sup>27</sup> An interesting application that has found wide use is the measurement of treatment effects and program effectiveness.<sup>28</sup>

An earnings equation that accounts for the value of a college education is

$$\text{earnings}_i = x_i' \beta + \delta C_i + \varepsilon_i,$$

where  $C_i$  is a dummy variable indicating whether or not the individual attended college. The same format has been used in any number of other analyses of programs, experiments, and treatments. The question is: Does  $\delta$  measure the value of a college education (assuming that the rest of the regression model is correctly specified)? The answer is

<sup>27</sup>For a survey, see Maddala (1983).

<sup>28</sup>This is one of the fundamental applications of this body of techniques, and is also the setting for the most longstanding and contentious debate on the subject. A *Journal of Business and Economic Statistics* symposium [Angrist (2001)] raised many of the important questions on whether and how it is possible to measure treatment effects.

19-60

TABLE 24.5 Empirical Distribution of Propensity Scores

Percent	Lower	Upper			
0-5	0.000591	0.000783		Sample size = 1,347	
5-10	0.000787	0.001061		Average score = 0.137238	
10-15	0.001065	0.001377		Std. Dev score = 0.274079	
15-20	0.001378	0.001748			
20-25	0.001760	0.002321			
25-30	0.002340	0.002956	1	Lower 0.000591	Upper 0.098016
30-35	0.002974	0.004057	2	0.098016	0.195440
35-40	0.004059	0.005272	3	0.195440	0.390289
40-45	0.005278	0.007486	4	0.390289	0.585138
45-50	0.007557	0.010451	5	0.585138	0.779986
50-55	0.010563	0.014643	6	0.779986	0.877411
55-60	0.014686	0.022462	7	0.877411	0.926123
60-65	0.022621	0.035060	8	0.926123	0.974835
65-70	0.035075	0.051415			
70-75	0.051415	0.076188			
75-80	0.076376	0.134189			
80-85	0.134238	0.320638			
85-90	0.321233	0.616002			
90-95	0.624407	0.949418			
95-100	0.949418	0.974835			

and (24-23) using the entire sample. The estimates of  $\delta$ ,  $\rho$ , and  $\sigma$  were (-1.01437, 0.35519, 1.38426). Using the results from the probit model, we averaged the result in (24-24) for the entire sample, obtaining an estimated treatment effect of \$1,476.30.

19.5.4 ~~24.5.7~~ SAMPLE SELECTION IN NONLINEAR MODELS

The preceding analysis has focused on an extension of the linear regression (or the estimation of simple averages of the data). The method of analysis changes in nonlinear models. To begin, it is not necessarily obvious what the impact of the sample selection is on the response variable, or how it can be accommodated in a model. Consider the model analyzed by Boyes, Hoffman, and Lowe (1989):

$$y_{i1} = 1 \text{ if individual } i \text{ defaults on a loan, 0 otherwise.}$$

$$y_{i2} = 1 \text{ if the individual is granted a loan, 0 otherwise.}$$

Wynand and van Praag (1981) also used this framework to analyze consumer insurance purchases in the first application of the selection methodology in a nonlinear model. Greene (1992) applied the same model to  $y_1$  = default on credit card loans, in which  $y_2$  denotes whether an application for the card was accepted or not. [Mohanty (2002) also used this model to analyze teen employment in California.] For a given individual,  $y_1$  is not observed unless  $y_2 = 1$ . Following the lead of the linear regression case in Section 24.5.3, a natural approach might seem to be to fit the second (selection) equation using a univariate probit model, compute the inverse Mills ratio,  $\lambda_i$ , and add it to the first equation as an additional "control" variable to accommodate the selection effect. [This is the approach used by Wynand and van Praag (1981).] The problems with this control function approach are, first, it is unclear what in the model is being "controlled" and, second, assuming the first model is correct, the appropriate model conditioned

AV: Confirm x-ref to Sec 24.5.3 is OK

and Greene (1994)

[see Terza (2010) for discussion.]

19-61

896 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

on the sample selection is unlikely to contain an inverse Mills ratio anywhere in it. That result is specific to the linear model, where it arises as  $E[\varepsilon_i | \text{selection}]$ . What would seem to be the apparent counterpart for this probit model,

$$\text{Prob}(y_{i1} = 1 | \text{selection on } y_{i2} = 1) = \Phi(x'_{i1}\beta_1 + \theta\lambda_i),$$

is not, in fact, the appropriate conditional mean, or probability. For this particular application, the appropriate conditional probability (extending the bivariate probit model of Section 23.8) would be

17.5

$$\text{Prob}[y_{i1} = 1 | y_{i2} = 1] = \frac{\Phi_2(x'_{i1}\beta_1, x'_{i2}\beta_2, \rho)}{\Phi(x'_{i2}\beta_2)}$$

We would use this result to build up the likelihood function for the three observed outcomes, as follows: The three types of observations in the sample, with their unconditional probabilities, are

$$\begin{aligned} y_{i2} = 0: \text{Prob}(y_{i2} = 0 | x_{i1}, x_{i2}) &= 1 - \Phi(x'_{i2}\beta_2), & 19-25 \\ y_{i1} = 0, y_{i2} = 1: \text{Prob}(y_{i1} = 0, y_{i2} = 1 | x_{i1}, x_{i2}) &= \Phi_2(-x'_{i1}\beta_1, x'_{i2}\beta_2, -\rho), & (24-25) \\ y_{i1} = 1, y_{i2} = 1: \text{Prob}(y_{i1} = 1, y_{i2} = 1 | x_{i1}, x_{i2}) &= \Phi_2(x'_{i1}\beta_1, x'_{i2}\beta_2, \rho). \end{aligned}$$

4N  
32

The log-likelihood function is based on these probabilities. <sup>33</sup> An application appears in Section 17.5.6.

Example 24.11 Doctor Visits and Insurance

Continuing our analysis of the utilization of the German health care system, we observe that the data set ~~(see Example 17.5)~~ contains an indicator of whether the individual subscribes to the "Public" health insurance or not. Roughly 87 percent of the observations in the sample do. We might ask whether the selection on public insurance reveals any substantive difference in visits to the physician. We estimated a logit specification for this model in Example 23.4

19-25

17.4

Using (24-25) as the framework, we define  $y_{i2}$  to be presence of insurance and  $y_{i1}$  to be the binary variable defined to equal 1 if the individual makes at least one visit to the doctor in the survey year. <sup>19.9</sup>

The estimation results are given in Table 24.6. Based on these results, there does appear to be a very strong relationship. The coefficients do change somewhat in the conditional model. A Wald test for the presence of the selection effect against the null hypothesis that  $\rho$  equals zero produces a test statistic of  $(-7.188)^2 = 51.667$ , which is larger than the critical value of 3.84. Thus, the hypothesis is rejected. A likelihood ratio statistic is computed as the difference between the log-likelihood for the full model and the sum of the two separate log-likelihoods for the independent probit models when  $\rho$  equals zero. The result is

1B  
19.9

$$\lambda_{LR} = 2[-23969.58 - (-15536.39 + (-8471.508))] = 77.796$$

17.5.3

The hypothesis is rejected once again. Partial effects were computed using the results in Section 23.8.3.

The large correlation coefficient can be misleading. The estimated  $-0.9299$  does not state that the presence of insurance makes it much less likely to go to the doctor. This is the correlation among the unobserved factors in each equation. The factors that make it more likely to purchase insurance make it less likely to use a physician. To obtain a simple correlation between the two variables, we might use the tetrachoric correlation defined in Example 23.13. This would be computed by fitting a bivariate probit model for the two binary variables without any other variables. The estimated value is 0.120.

17.18

32  
33

Extensions of the bivariate probit model to other types of censoring are discussed in Poirier (1980) and Abowd and Farber (1982).

19-62

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 897

19.9  
TABLE 24.6 Estimated Probit Equations for Doctor Visits

Variable	Independent: No Selection			Sample Selection Model		
	Estimate	Standard Error	Partial Effect	Estimate	Standard Error	Partial Effect
Constant	0.05588	0.06564		-9.4366	0.06760	
Age	0.01331	0.0008399	0.004971	0.01284	0.0008131	0.005042
Income	-0.1034	0.05089	-0.03860	-0.1030	0.04582	-0.04060
Kids	-0.1349	0.01947	-0.05059	-0.1264	0.01790	-0.04979
Education	-0.01920	0.004254	-0.007170	0.03660	0.004744	0.002703
Married	0.03586	0.02172	0.01343	0.03564	0.02016	0.01404
ln L	-15536.39					
Constant	3.3585	0.06959		3.2699	0.06916	
Age	0.0001868	0.0009744		-0.0002679	0.001036	
Education	-0.1854	0.003941		-0.1807	0.003936	
Female	0.1150	0.02186	0.0000 <sup>a</sup>	0.2230	0.02101	0.01446 <sup>a</sup>
ln L	-8471.508					
$\rho$	0.0000	0.0000		-0.9299	0.1294	
ln L	-24007.90			-23969.58		

<sup>a</sup>Indirect effect from second equation.

✓  
✓  
✓

present

2005 → 2010

1998, 2010

More general cases are typically much less straightforward. Greene (1992, 2006) and Terza (1995, 1998) have developed sample selection models for nonlinear specifications based on the underlying logic of the Heckman model in Section 24.5.3; that the influence of the incidental truncation acts on the unobservable variables in the model. (That is the source of the "selection bias" in conventional estimators.) The modeling extension introduces the unobservables into the model in a natural fashion that parallels the regression model. Terza (2010) presents a survey of the general results.

19.5.3

The generic model will take the form

1. Probit selection equation:

$$z_i^* = \mathbf{w}_i' \boldsymbol{\alpha} + u_i \text{ in which } u_i \sim N[0, 1],$$

$$z_i = 1 \text{ if } z_i^* > 0, 0 \text{ otherwise.}$$

19 (24.26)

2. Nonlinear index function model with unobserved heterogeneity and sample selection:

$$\mu_i | \varepsilon_i = \mathbf{x}_i' \boldsymbol{\beta} + \sigma \varepsilon_i, \varepsilon_i \sim N[0, 1],$$

$$y_i | \mathbf{x}_i, \varepsilon_i \sim \text{density } g(y_i | \mathbf{x}_i, \varepsilon_i) = f(y_i | \mathbf{x}_i' \boldsymbol{\beta} + \sigma \varepsilon_i),$$

$$y_i, \mathbf{x}_i \text{ are observed only when } z_i = 1,$$

$$[u_i, \varepsilon_i] \sim N[(0, 1), (1, \rho, 1)].$$

19 (24.27)

used

For example, in a Poisson regression model, the conditional mean function becomes  $E(y_i | \mathbf{x}_i) = \lambda_i = \exp(\mathbf{x}_i' \boldsymbol{\beta} + \sigma \varepsilon_i) = \exp(\mu_i)$ . (We will use this specification of the model in Chapter 25 to introduce random effects in the Poisson regression model for panel data.)

The log-likelihood function for the full model is the joint density for the observed data. When  $z_i$  equals one,  $(y_i, \mathbf{x}_i, z_i, \mathbf{w}_i)$  are all observed. To obtain the joint density

Handwritten note: This term "heterogeneity" was KT on msp 19-116. Here also?

19-63

898 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

$p(y_i, z_i = 1 | x_i, w_i)$ , we proceed as follows:

$$p(y_i, z_i = 1 | x_i, w_i) = \int_{-\infty}^{\infty} p(y_i, z_i = 1 | x_i, w_i, \varepsilon_i) f(\varepsilon_i) d\varepsilon_i.$$

Conditioned on  $\varepsilon_i$ ,  $z_i$  and  $y_i$  are independent. Therefore, the joint density is the product,

$$p(y_i, z_i = 1 | x_i, w_i, \varepsilon_i) = f(y_i | x_i' \beta + \sigma \varepsilon_i) \text{Prob}(z_i = 1 | w_i, \varepsilon_i).$$

The first part,  $f(y_i | x_i' \beta + \sigma \varepsilon_i)$  is the conditional index function model in (24-27). By joint normality,  $f(u_i | \varepsilon_i) = N[\rho \varepsilon_i, (1 - \rho^2)]$ , so  $u_i | \varepsilon_i = \rho \varepsilon_i + (u_i - \rho \varepsilon_i) = \rho \varepsilon_i + v_i$  where  $E[v_i] = 0$  and  $\text{Var}[v_i] = (1 - \rho^2)$ . Therefore,

$$\text{Prob}(z_i = 1 | w_i, \varepsilon_i) = \Phi \left( \frac{w_i' \alpha + \rho \varepsilon_i}{\sqrt{1 - \rho^2}} \right).$$

Combining terms and using the earlier approach, the unconditional joint density is

$$p(y_i, z_i = 1 | x_i, w_i) = \int_{-\infty}^{\infty} f(y_i | x_i' \beta + \sigma \varepsilon_i) \Phi \left( \frac{w_i' \alpha + \rho \varepsilon_i}{\sqrt{1 - \rho^2}} \right) \frac{\exp(-\varepsilon_i^2/2)}{\sqrt{2\pi}} d\varepsilon_i. \quad (24-28)$$

The other part of the likelihood function for the observations with  $z_i = 0$  will be

$$\begin{aligned} \text{Prob}(z_i = 0 | w_i) &= \int_{-\infty}^{\infty} \text{Prob}(z_i = 0 | w_i, \varepsilon_i) f(\varepsilon_i) d\varepsilon_i. \\ &= \int_{-\infty}^{\infty} \left[ 1 - \Phi \left( \frac{w_i' \alpha + \rho \varepsilon_i}{\sqrt{1 - \rho^2}} \right) \right] f(\varepsilon_i) d\varepsilon_i. \\ &= \int_{-\infty}^{\infty} \Phi \left( \frac{-(w_i' \alpha + \rho \varepsilon_i)}{\sqrt{1 - \rho^2}} \right) \frac{\exp(-\varepsilon_i^2/2)}{\sqrt{2\pi}} d\varepsilon_i. \end{aligned} \quad (24-29)$$

For convenience, we can use the invariance principle to reparameterize the likelihood function in terms of  $\gamma = \alpha / \sqrt{1 - \rho^2}$  and  $\tau = \rho / \sqrt{1 - \rho^2}$ . Combining all the preceding terms, the log-likelihood function to be maximized is

$$\ln L = \sum_{i=1}^n \ln \int_{-\infty}^{\infty} [(1 - z_i) + z_i f(y_i | x_i' \beta + \sigma \varepsilon_i)] \Phi[(2z_i - 1)(w_i' \gamma + \tau \varepsilon_i)] \phi(\varepsilon_i) d\varepsilon_i. \quad (24-30)$$

This can be maximized with respect to  $(\beta, \sigma, \gamma, \tau)$  using quadrature or simulation. When done,  $\rho$  can be recovered from  $\rho = \frac{\tau}{\sqrt{1 + \tau^2}}$  and  $\alpha = (1 - \rho^2)^{1/2} \gamma$ . All that differs from one model to another is the specification of  $f(y_i | x_i' \beta + \sigma \varepsilon_i)$ . This is the specification used in Greene (1992), Terza (1998), and Terza and Kenkel (2001). (In these two papers, the authors analyzed  $E[y_i | z_i = 1]$  rather than the conditional density. Their estimator was based on nonlinear least squares, but as earlier, it is necessary to integrate the unobserved heterogeneity out of the conditional mean function.)

also  
this

Greene (2010) applies the method to a stochastic frontier model.

19.5.5 ~~24.5.8~~ PANEL DATA APPLICATIONS OF SAMPLE SELECTION MODELS

The development of methods for extending sample selection models to panel data settings parallels the literature on cross-section methods. It begins with Hausman and Wise (1979) who devised a maximum likelihood estimator for a two-period model with attrition—the “selection equation” was a formal model for attrition from the sample.

19-64

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 899

Subsequent research has drawn the analogy between attrition and sample selection in a variety of applications, such as Keane et al. (1988) and Verbeek and Nijman (1992), and produced theoretical developments including Wooldridge (2002a, b).

The direct extension of panel data methods to sample selection brings several new issues for the modeler. An immediate question arises concerning the nature of the selection itself. Although much of the theoretical literature [e.g., Kyriazidou (1997, 2001)] treats the panel as if the selection mechanism is run anew in every period, in practice, the selection process often comes in two very different forms. First, selection may take the form of selection of the entire group of observations into the panel data set. Thus, the selection mechanism operates once, perhaps even before the observation window opens. Consider the entry (or not) of eligible candidates for a job training program. In this case, it is not appropriate to build the model to allow entry, exit, and reentry. Second, for most applications, selection comes in the form of attrition or retention. Once an observation is "deselected," it does not return. Leading examples would include "survivorship" in time-series-cross-section models of firm performance and attrition in medical trials and in panel data applications involving large national survey data bases, such as Contoyannis et al. (2004). Each of these cases suggests the utility of a more structured approach to the selection mechanism.

19.5.5. a ~~24.5.5. a~~ Common Effects in Sample Selection Models

A formal "effects" treatment for sample selection was first suggested in complete form by Verbeek (1990), who formulated a random effects model for the probit equation and a fixed effects approach for the main regression. Zabel (1992) criticized the specification for its asymmetry in the treatment of the effects in the two equations. He also argued that the likelihood function that neglected correlation between the effects and regressors in the probit model would render the FIML estimator inconsistent. His proposal involved fixed effects in both equations. Recognizing the difficulty of fitting such a model, he then proposed using the Mundlak correction. The full model is

$$\begin{aligned} y_{it}^* &= \eta_i + \mathbf{x}'_{it}\beta + \varepsilon_{it}, & \eta_i &= \bar{\mathbf{x}}'_i\pi + \tau w_i, & w_i &\sim N[0, 1], \\ d_{it}^* &= \theta_i + \mathbf{z}'_{it}\alpha + u_{it}, & \theta_i &= \bar{\mathbf{z}}'_i\delta + \omega v_i, & v_i &\sim N[0, 1], \\ (\varepsilon_{it}, u_{it}) &\sim N_2[(0, 0), (\sigma^2, 1, \rho\sigma)]. \end{aligned} \tag{24-31}$$

The "selectivity" in the model is carried through the correlation between  $\varepsilon_{it}$  and  $u_{it}$ . The resulting log-likelihood is built up from the contribution of individual  $i$ .

$$\begin{aligned} L_i &= \int_{-\infty}^{\infty} \prod_{d_{it}=0} \Phi[-\mathbf{z}'_{it}\alpha - \bar{\mathbf{z}}'_i\delta - \omega v_i] \phi(v_i) dv_i \\ &\times \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \prod_{d_{it}=1} \Phi \left[ \frac{\mathbf{z}'_{it}\alpha + \bar{\mathbf{z}}'_i\delta + \omega v_i + (\rho/\sigma)\varepsilon_{it}}{\sqrt{1-\rho^2}} \right] \\ &\times \frac{1}{\sigma} \phi\left(\frac{\varepsilon_{it}}{\sigma}\right) \phi_2(v_i, w_i) dv_i dw_i. \end{aligned} \tag{24-32}$$

$$\varepsilon_{it} = y_{it} - \mathbf{x}'_{it}\beta - \bar{\mathbf{x}}'_i\pi - \tau w_i.$$

The log-likelihood is then  $\ln L = \sum_i \ln L_i$ .

The log-likelihood requires integration in two dimensions for any selected observations. Vella (1998) suggested two-step procedures to avoid the integration. However,

## 900 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

the bivariate normal integration is actually the product of two univariate normals, because in the preceding specification,  $v_i$  and  $w_i$  are assumed to be uncorrelated. As such, the likelihood function in (24-32) can be readily evaluated using familiar simulation or quadrature techniques. [See Sections 16.9.6.b and 47.5. Vella and Verbeek (1999) suggest this in a footnote, but do not pursue it.] To show this, note that the first line in the log-likelihood is of the form  $E_v[\prod_{d=0} \Phi(\dots)]$  and the second line is of the form  $E_{vw}[E_v[\Phi(\dots)\phi(\dots)/\sigma]]$ . Either of these expectations can be satisfactorily approximated with the average of a sufficient number of draws from the standard normal populations that generate  $w_i$  and  $v_i$ . The term in the simulated likelihood that follows this prescription is

$$L_i^S = \frac{1}{R} \sum_{r=1}^R \prod_{d_i=0} \Phi[-z'_{it}\alpha - \bar{z}'_i\delta - \omega v_{i,r}] \times \frac{1}{R} \sum_{r=1}^R \prod_{d_i=1} \Phi \left[ \frac{z'_{it}\alpha + \bar{z}'_i\delta + \omega v_{i,r} + (\rho/\sigma)\varepsilon_{it,r}}{\sqrt{1-\rho^2}} \right] \frac{1}{\sigma} \phi \left( \frac{\varepsilon_{it,r}}{\sigma} \right), \quad (24-33)$$

$$\varepsilon_{it,r} = y_{it} - x'_{it}\beta - \bar{x}'_i\pi - \tau w_{i,r}.$$

Maximization of this log-likelihood with respect to  $(\beta, \sigma, \rho, \alpha, \delta, \pi, \tau, \omega)$  by conventional gradient methods is quite feasible. Indeed, this formulation provides a means by which the likely correlation between  $v_i$  and  $w_i$  can be accommodated in the model. Suppose that  $w_i$  and  $v_i$  are bivariate standard normal with correlation  $\rho_{vw}$ . We can project  $w_i$  on  $v_i$  and write

$$w_i = \rho_{vw}v_i + (1 - \rho_{vw}^2)^{1/2}h_i,$$

where  $h_i$  has a standard normal distribution. To allow the correlation, we now simply substitute this expression for  $w_i$  in the simulated (or original) log-likelihood and add  $\rho_{vw}$  to the list of parameters to be estimated. The simulation is still over independent normal variates,  $v_i$  and  $h_i$ .

Notwithstanding the preceding derivation, much of the recent attention has focused on simpler two-step estimators. Building on Ridder and Wansbeek (1990) and Verbeek and Nijman (1992) [see Vella (1998) for numerous additional references], Vella and Verbeek (1999) propose a two-step methodology that involves a random effects framework similar to the one in (24-31). As they note, there is some loss in efficiency by not using the FIML estimator. But, with the sample sizes typical in contemporary panel data sets, that efficiency loss may not be large. As they note, their two-step template encompasses a variety of models including the tobit model examined in the preceding sections and the mover-stayer model noted earlier.

The Vella and Verbeek model requires some fairly intricate maximum likelihood procedures. Wooldridge (1995) proposes an estimator that, with a few probably—but not necessarily—innocent assumptions, can be based on straightforward applications of conventional, everyday methods. We depart from a fixed effects specification,

$$y_{it}^* = \eta_i + x'_{it}\beta + \varepsilon_{it},$$

$$d_{it}^* = \theta_i + x'_{it}\alpha + u_{it},$$

$$(\varepsilon_{it}, u_{it}) \sim N_2[(0, 0), (\sigma^2, 1, \rho\sigma)].$$

CHAPTER 24 ♦ Truncation, Censoring, and Sample Selection 901

Under the **mean independence assumption**  $E[\varepsilon_{it} | \eta_i, \theta_i, z_{i1}, \dots, z_{iT}, v_{i1}, \dots, v_{iT}, d_{i1}, \dots, d_{iT}] = \rho u_{it}$ , it will follow that

$$E[y_{it} | x_{i1}, \dots, x_{iT}, \eta_i, \theta_i, z_{i1}, \dots, z_{iT}, v_{i1}, \dots, v_{iT}, d_{i1}, \dots, d_{iT}] = \eta_i + x'_{it}\beta + \rho u_{it}.$$

This suggests an approach to estimating the model parameters; however, it requires computation of  $u_{it}$ . That would require estimation of  $\theta_i$ , which cannot be done, at least not consistently—and that precludes simple estimation of  $u_{it}$ . To escape the dilemma, Wooldridge (2002c) suggests Chamberlain's approach to the fixed effects model,

$$\theta_i = f_0 + z'_{i1}f_1 + z'_{i2}f_2 + \dots + z'_{iT}f_T + h_i.$$

With this substitution,

$$\begin{aligned} d_{it}^* &= z'_{it}\alpha + f_0 + z'_{i1}f_1 + z'_{i2}f_2 + \dots + z'_{iT}f_T + h_i + u_{it} \\ &= z'_{it}\alpha + f_0 + z'_{i1}f_1 + z'_{i2}f_2 + \dots + z'_{iT}f_T + w_{it}, \end{aligned}$$

where  $w_{it}$  is independent of  $z_{it}$ ,  $t = 1, \dots, T$ . This now implies that

$$\begin{aligned} E[y_{it} | x_{i1}, \dots, x_{iT}, \eta_i, \theta_i, z_{i1}, \dots, z_{iT}, v_{i1}, \dots, v_{iT}, d_{i1}, \dots, d_{iT}] &= \eta_i + x'_{it}\beta + \rho(w_{it} - h_i) \\ &= (\eta_i - \rho h_i) + x'_{it}\beta + \rho w_{it}. \end{aligned}$$

To complete the estimation procedure, we now compute  $T$  cross-sectional probit models (reestimating  $f_0, f_1, \dots$  each time) and compute  $\hat{\lambda}_{it}$  from each one. The resulting equation,

$$y_{it} = a_i + x'_{it}\beta + \rho \hat{\lambda}_{it} + v_{it},$$

now forms the basis for estimation of  $\beta$  and  $\rho$  by using a conventional fixed effects linear regression with the observed data.

19.5.5.b

~~24.5.8.6~~ Attrition

The recent literature on sample selection contains numerous analyses of two-period models, such as Kyriazidou (1997, 2001). They generally focus on non- and semiparametric analyses. An early parametric contribution of Hausman and Wise (1979) is also a two-period model of attrition, which would seem to characterize many of the studies suggested in the current literature. The model formulation is a two-period random effects specification:

$$y_{i1} = x'_{i1}\beta + \varepsilon_{i1} + u_i \quad (\text{first period regression}),$$

$$y_{i2} = x'_{i2}\beta + \varepsilon_{i2} + u_i \quad (\text{second period regression}).$$

Attrition is likely in the second period (to begin the study, the individual must have been observed in the first period). The authors suggest that the probability that an observation is made in the second period varies with the value of  $y_{i2}$  as well as some other variables.

$$z_{i2}^* = \delta y_{i2} + x'_{i2}\theta + w'_{i2}\alpha + v_{i2}.$$

Attrition occurs if  $z_{i2}^* \leq 0$ , which produces a probit model,

$$z_{i2} = 1 (z_{i2}^* > 0) \quad (\text{attrition indicator observed in period 2}).$$

## 902 PART VI ♦ Cross Sections, Panel Data, and Microeconometrics

An observation is made in the second period if  $z_{i2} = 1$ , which makes this an early version of the familiar sample selection model. The reduced form of the observation equation is

$$\begin{aligned} z_{i2}^* &= \mathbf{x}'_{i2}(\delta\beta + \theta) + \mathbf{w}'_{i2}\alpha + \delta\varepsilon_{i2} + v_{i2} \\ &= \mathbf{x}'_{i2}\pi + \mathbf{w}'_{i2}\alpha + h_{i2} \\ &= \mathbf{x}'_{i2}\gamma + h_{i2}. \end{aligned}$$

The variables in the probit equation are all those in the second period regression plus any additional ones dictated by the application. The estimable parameters in this model are  $\beta, \gamma, \sigma^2 = \text{Var}[\varepsilon_{it} + u_i]$ , and two correlation coefficients,

$$\rho_{12} = \text{Corr}[\varepsilon_{i1} + u_i, \varepsilon_{i2} + u_i] = \text{Var}[u_i]/\sigma^2,$$

and

$$\rho_{23} = \text{Corr}[h_{i2}, \varepsilon_{i2} + u_i].$$

All disturbances are assumed to be normally distributed. (Readers are referred to the paper for motivation and details on this specification.)

The authors propose a full information maximum likelihood estimator. Estimation can be simplified somewhat by using two steps. The parameters of the probit model can be estimated first by maximum likelihood. Then the remaining parameters are estimated by maximum likelihood, conditionally on these first-step estimates. The Murphy and Topel adjustment is made after the second step. [See Greene (2007a).]

The Hausman and Wise model covers the case of two periods in which there is a formal mechanism in the model for retention in the second period. It is unclear how the procedure could be extended to a multiple period application such as that in Contoyannis et al. (2004), which involved a panel data set with eight waves. In addition, in that study, the variables in the main equations were counts of hospital visits and physician visits, which complicates the use of linear regression. A workable solution to the problem of attrition in a multiperiod panel is the **inverse probability weighted estimator** [Wooldridge (2002a, 2006b) and Rotnitzky and Robins (2005).] In the Contoyannis application, there are eight waves in the panel. Attrition is taken to be "ignorable" so that the unobservables in the attrition equation and in the main equation(s) of interest are uncorrelated. (Note that Hausman and Wise do not make this assumption.) This enables Contoyannis et al. to fit a "retention" probit equation for each observation present at wave 1, for waves 2-8, using characteristics observed at the entry to the panel. (This defines, then, "selection (retention) on observables.") Defining  $d_{it}$  to be the indicator for presence ( $d_{it} = 1$ ) or absence ( $d_{it} = 0$ ) of observation  $i$  in wave  $t$ , it will follow that the sequence of observations will begin at 1 and either stay at 1 or change to 0 for the remaining waves. Let  $\hat{p}_{it}$  denote the predicted probability from the probit estimator at wave  $t$ . Then, their full log-likelihood is constructed as

$$\ln L = \sum_{i=1}^n \sum_{t=1}^T \frac{d_{it}}{\hat{p}_{it}} \ln L_{it}.$$

Wooldridge (2002b) presents the underlying theory for the properties of this weighted maximum likelihood estimator. [Further details on the use of the inverse probability weighted estimator in the Contoyannis et al. (2004) study appear in Jones, Koolman, and Rice (2006) and in Section 17.4.9.]

## 19.6 Evaluating Treatment Effects

The leading recent application of models of selection and endogeneity is the evaluation of "treatment effects." The central focus is on analysis of the effect of participation in a treatment,  $T$ , on an outcome variable,  $y$  - examples include job training programs [LaLonde (1986), Business Week (2009, below)] and education [e.g., test scores, Angrist and Lavy (1999), Van der Klaauw (2002)]. Wooldridge and Imbens (2009, pp. 22-23) cite a number of labor market applications. Recent more narrow examples include Munkin and Trivedi's (2007) analysis of the effect of dental insurance and Jones and Rice's (2010) survey that notes a variety of techniques and applications in health economics.

Example 19.14

### Example 19.14 German Labor Market Interventions

"Germany long had the highest ratio of unfilled jobs to unemployed people in Europe. Then, in 2003, Berlin launched the so-called Hartz reforms, ending generous unemployment benefits that went on indefinitely. Now payouts for most recipients drop sharply after a year, spurring people to look for work. From 12.7% in 2005, unemployment fell to 7.1% last November. Even now, after a year of recession, Germany's jobless rate has risen to just 8.6%.

At the same time, lawmakers introduced various programs intended to make it easier for people to learn new skills. One initiative instructed the Federal Labor Agency, which had traditionally pushed the long-term unemployed into government-funded make-work positions, to cooperate more closely with private employers to create jobs. That program last year paid Dutch staffing agency Randstad to teach 15,000 Germans information technology, business English, and other skills. And at a Daimler truck factory in Wörth, 55 miles west of Stuttgart, several dozen short-term employees at risk of being laid off got government help to continue working for the company as mechanic trainees.

umlauf

Under a second initiative, Berlin pays part of the wages of workers hired from the ranks of the jobless. Such payments make employers more willing to take on the costs of training new workers. That extra training, in turn, helps those workers keep their jobs after the aid expires, a study by the government-funded Institute for Employment Research found. Café Nenninger in the city of Kassel, for instance, used the program to train an unemployed single mother. Co-owner Verena Nenninger says she was willing to take a chance on her in part because the government picked up about a third of her salary the first year. "It was very helpful, because you never know what's going to happen," Nenninger says. [Business Week (2009).]

acute

note single quotes

Empirical measurement of treatment effects, such as the impact of going to college or participating in a job training program, presents a large variety of econometric complications. The natural, ultimate objective of an analysis of a "treatment" or intervention would be the "effect of treatment on the treated." For example, what is the effect of a college education on the lifetime income of someone who goes to college? Measuring this effect econometrically encounters at least two compelling computations:

**Endogeneity of the treatment:** The analyst risks attributing to the treatment causal effects that should be attributed to factors that motivate both the treatment and the outcome. In our example, the individual who goes to college might well have succeeded (more) in life than their counterpart who did not go to college even if they (themselves) did not attend college.

**Missing counterfactual:** The preceding thought experiment is not actually the effect we wish to measure. In order to measure the impact of college attendance on lifetime earnings in a pure sense, we would have to run an individual's lifetime twice, once with college attendance and once without. Any individual is observed in only one of the two states, so the pure measurement is impossible.

Accommodating these two problems forms the focal point of this enormous and still growing literature. **Rubin's causal model** (1974, 1978) provides a useful framework for the analysis. Every individual in a population has a potential outcome,  $y$  and can be exposed to the treatment,  $C$ . We will denote by  $C_i$  the indicator whether or not the individual receives the treatment. Thus, the potential outcomes are  $y_i(C_i = 1) = y_{i1}$  and  $y_i(C_i = 0) = y_{i0}$ . The **average treatment effect**, averaged across the entire population is

$$ATE = E[y_{i1} - y_{i0}].$$

The compelling complication is that the individual will exist in only one of the two states, so it is not possible to estimate ATE without further assumptions. More specifically, what the researcher would prefer see is the **average treatment effect on the treated**.

$$ATET = E[y_{i1} - y_{i0} | C_i = 1]$$

and note that the second term is the missing counterfactual.

One of the major themes of the recent research is to devise robust methods of estimation that do not rely heavily on fragile assumptions such as identification by functional form (e.g., relying on bivariate normality) and identification by exclusion restrictions (e.g., relying on basic instrumental variable estimators). This is a challenging exercise — we have relied heavily on these assumptions in most of the work in this book up to this point. For purposes of the general specification, we will denote by  $\mathbf{x}$ , the exogenous information that will be brought to bear on this estimation problem. The vector  $\mathbf{x}$  may (usually will) be a set of variables that will appear in a regression model, but it is useful to think more generally than that, and consider  $\mathbf{x}$  rather to be an information set. Certain minimal assumptions are necessary to make any headway at all. The following appear at different points in the analysis.

**Conditional independence:** Receiving the treatment,  $C_i$ , does not depend on the outcome variable once the effect of  $\mathbf{x}$  on the outcome is accounted for. If assignment to the treatment group is completely random, then we would omit the effect of  $\mathbf{x}$  in this assumption. This assumption is extended for regression approaches with the **conditional mean assumption:**  $E[y_{i0} | \mathbf{x}_i, C_i = 1] = E[y_{i0} | \mathbf{x}_i, C_i = 0] = E[y_{i0} | \mathbf{x}]$ . This states that the outcome in the untreated state does not affect the participation.

**Distribution of potential outcomes:** The model that is used for the outcomes is the same for treated and nontreated,  $f(y | \mathbf{x}, T=1) = f(y | \mathbf{x}, T=0)$ . In a regression context, this would mean that the same regression applies in both states and that the disturbance is uncorrelated with  $T$ , or that  $T$  is exogenous. This is a very strong assumption that we will relax below. For the present, it removes one of the complications noted above, so a step in the model building exercise will be to relax this assumption.

**Overlap assumption:** For any value of  $\mathbf{x}$ ,  $0 < \text{Prob}(C_i = 1 | \mathbf{x}) < 1$ . The strict inequality in this assumption means that for any  $\mathbf{x}$ , the population will contain a mix of treated and nontreated individuals. The usefulness of the overlap assumption is that with it, we can expect to find, for any treated individual, an individual who looks like them, but is not treated. This assumption will be useful for regression approaches.

that/



KT  
 Au: This KT  
 term is  
 not in  
 chap. list

later/

previously/

sections

The following will describe three major parts of the research agenda on treatment effects: regression analysis with control functions in Section 19.6.1, propensity score matching in Section 19.6.2, and regression discontinuity design in Section 19.6.3. A fourth area, instrumental variable estimation, was developed in Chapter 8. As noted, this is a huge and rapidly growing literature. For example, Imbens and Wooldridge's (2009) survey paper runs to 85 pages, includes nearly 300 references, most of them since 2000. Our purpose here is to provide some of the vocabulary and a superficial introduction to methods. The survey papers by Imbens and Wooldridge (2009) and Jones and Rice (2010) provide greater detail. The conference volume by Millment, Smith, and Vytlačil (2008) contains many theoretical contributions and empirical applications.<sup>34</sup> A *Journal of Business and Economic Statistics* symposium [Angrist (2001)] raised many of the important questions on whether and how it is possible to measure treatment effects.

✓  
✓  
✓

FU  
33

33/

33 In

<sup>34</sup> The initial essay in the volume, Goldberger (2008) reproduces Goldberger (1972) in which the author explores the endogeneity issue in detail with specific reference to the Head Start program of the 1960s.

✓✓